

611-TD-598-001

EOSDIS Core System Project

M&O Procedures: Section 17—Archive Procedures

Interim Update

January 2002

Raytheon Systems Company
Upper Marlboro, Maryland

Preface

This document is an interim update to the Mission Operations Procedures Manual for the ECS Project, document number 611-CD-600-001. This document has not been submitted to NASA for approval, and should be considered unofficial.

This update has been prepared to improve the logical organization and format of the Archive Procedures and to incorporate additional revisions appropriate for Release 6A and Synergy. It constitutes a complete rewriting of the Archive Procedures section.

Any questions should be addressed to:

Data Management Office
The ECS Project Office
Raytheon Systems Company
1616 McCormick Drive
Upper Marlboro, Maryland 20774-5301

This page intentionally left blank.

Contents

17. ARCHIVE PROCEDURES	1
17.1 Starting and Stopping AMASS	2
17.1.1 Starting the AMASS Application	2
17.1.2 Shutting Down AMASS Tape Archive System	4
17.1.3 Rebooting AMASS	4
17.2 Loading, Removing, and Managing Archive Media	5
17.2.1 Automatically Loading Archive Media	6
17.2.2 Manually Loading Archive Media	7
17.2.3 Formatting a Volume	9
17.2.4 Removing Media	10
17.3 Monitoring and Managing the Archive with ECS Custom GUIs	11
17.3.1 Launching DSS GUIs	11
17.3.2 Using Storage Management GUIs to Display Archive Path Information	13
17.3.3 Monitoring Archive Requests using the Storage Management GUI	15
17.3.4 Monitoring Distribution Requests using the Data Distribution GUI	17
17.3.5 Setting Checksum Calculation	19
17.4 Deleting Granules	21
17.4.1 Resetting the Lock on the DsMdDeletedGranules Table	23
17.4.2 Selecting Granules for Deletion	24
17.4.2.1 Selection by ESDT ShortName, Version, and Granule Time Coverage	24
17.4.2.2 Selection by ESDT ShortName, Version, and Granule Insert Time Range	26
17.4.2.3 Selection Using a Separate Input File	28
17.4.3 Deleting Granules from the Inventory and Archive (Physical Deletion)	30
17.4.4 Deleting Granules from the Archive	32
17.5 Backing Up and Restoring AMASS	34
17.5.1 Creating a Backup for AMASS	34
17.5.2 Replacing the AMASS Database Backup Volume (Volume 1)	35
17.5.3 Restoring the AMASS Database	37
17.6 Backing Up and Restoring Archived Data	38
17.6.1 Backups for Archive Data	39
17.6.1.1 Creating Offsite Backups	39
17.6.1.2 Creating Replacement Backups Manually from Existing Archives	42
17.6.2 Restoring Archive Data	43
17.6.2.1 Manual Data Recovery from Local Backup Tapes	44
17.6.2.2 Manual Data Recovery from Offsite Backup Tapes	45
17.6.2.3 Manual Data Recovery from Damaged Cartridge	48
17.6.2.4 Data Recovery for Known Files Not Backed Up in ECS	52
17.6.2.4.1 Re-Ingest of Lost Data	53
17.6.2.4.2 Recovery of Lost Data by Reprocessing	53
17.6.2.4.2 Recovering Granules from Another Producing Site	54
17.6.2.5 SDSRV Procedures in Support of Data Recovery	55
17.6.2.5.1 SDSRV Retrieval of File Location Metadata	55
17.6.2.5.2 SDSRV Retrieval of Granule Production History Metadata	58

17.7 Archive Troubleshooting.....	62
17.7.1 Using AMASS Commands, Utilities, and Scripts for Monitoring and Fault Response.....	63
17.7.1.1 Checking Daemons and Using <i>healthcheck</i>	64
17.7.1.2 Using <i>sysperf</i> to Display the Status of AMASS I/O Activity.....	67
17.7.1.3 Using <i>vollist</i> to Display Volume Data.....	68
17.7.1.4 Using the <i>amass_log</i> Script to Display AMASS Errors.....	69
17.7.1.5 Using <i>quedisplay</i> to View What is in the AMASS Queue.....	70
17.7.1.6 Using <i>mediamove</i> to Establish Synchrony Between <i>quedisplay</i> and <i>medialist</i>	71
17.7.2 Recovering from Failure to Store or Retrieve Data.....	72
17.7.2.1 Checking Server Log Files.....	72
17.7.2.2 A Special Case: Checking the Request Manager Server Debug Log.....	79
17.7.2.3 Checking the <i>el_ETAC</i> Log.....	91
17.7.2.4 Handling a Data Insertion Failure.....	94
17.7.2.5 Handling a Data Acquire Failure.....	99
17.7.3 Diagnosing/Investigating Write Errors.....	104
17.7.4 Diagnosing/Investigating Read Errors.....	106
17.8 ACSLS Procedures.....	108
17.8.1 Entering the Archive After AMASS is Started.....	110
17.8.2 Backing Up the ACSLS Database.....	111
17.8.3 Restoring the ACSLS Database.....	112
17.8.4 Checking Cleaning Cartridges.....	113
17.9 Using the AMASS Graphical User Interface (GUI).....	114
17.9.1 Launching the AMASS GUI and Viewing Volume Group and Volume Information.....	115
17.9.2 Using the AMASS GUI to Modify a Volume Group.....	115
17.9.3 Modify a Volume.....	118
17.10 Data Pool Maintenance Tasks.....	121
17.10.1 Launch the Data Pool Maintenance (DPM) GUI.....	122
17.10.2 Using the DPM GUI to Monitor Data Pool Active Insert Processes.....	124
17.10.3 Suspend and Resume Data Pool Insert Actions.....	125
17.10.4 Check the Data Pool Insert Queue and Cancel a Data Pool Insert Action.....	126
17.10.5 Toggle the State of the NoFreeSpaceFlag.....	127
17.10.6 Configure the Number of Allowed Active Insert Processes.....	128
17.10.7 Configure the Default Retention Period and the Default Retention Priority.....	129
17.10.8 Use the DPM GUI to View Collection Groups and Collections.....	131
17.10.9 Use the DPM GUI to Modify Collection Groups.....	133
17.10.10 Use the DPM GUI to Add a Collection Group.....	134
17.10.11 Use the DPM GUI to Add Collections to a Collection Group.....	136
17.10.12 Use the DPM GUI to Modify Collections in a Collection Group.....	138
17.10.13 Use the Update Granule Utility to Extend the Retention for Selected Science Granules.....	140
17.10.14 Invoke the Data Pool Cleanup Utility Manually.....	144
17.10.15 Establish Data Pool Cleanup to Run (at 1:00 am) with <i>cron</i>	146
17.10.16 Specify Data Pool Access Statistics Rollup Start Time (at 1:00 am) and DPASU Execution (at 2:00 am), OPS Mode, with <i>cron</i>	148
17.10.17 Specify Data Pool Access Statistics Utility Execution from the Command Line.....	153
17.10.18 Archive Access Statistics using the Data Pool Archive Access Statistics Data Utility.....	156
17.10.19 Delete Access Statistics using the Data Pool Delete Access Statistics Data Utility.....	157
17.10.20 Restore Access Statistics using the Data Pool Restore Access Statistics Data Utility.....	158

17. Archive Procedures

Archive processing procedures support and maintain the process by which the Data Server Subsystem (DSS) manages persistent storage of earth science and related data, and through which the DSS provides search and retrieval access to the data. Through archive processing, data products that have been ingested into the system or produced by data processing on previously stored data are archived to tape for permanent storage and distributed to users via hard media (tape or disk) or electronic means. The DAAC Archive Manager's job entails working with the Science Data Specialist, the Science Coordinator, and the Resource Manager, as well as providing direction for the Data Ingest Technician. The physical archive is one or more StorageTek (STK) Powderhorn Model 9310 Automated Cartridge System tape storage towers, providing a mass storage system of jukeboxes for removable media (tape cartridges). The File Storage Management System (FSMS) software, hosted on a Silicon Graphics Inc. (SGI) Challenge XL or on an SGI Origin 2000, is the Archival Management and Storage System (AMASS), a product of Advanced Digital Information Corporation (ADIC). AMASS is a UNIX file system that manages files, volumes (media), drives and jukeboxes. The *AMASS System Administrator's User Guide* can be viewed using Adobe Acrobat and is available electronically on **drg** servers (e.g., g0drg01, e0drg11, l0drg01, n0drg01) in directory **/usr/amass/books**.

Archive processing activities include operating functions associated with the AMASS software, managing and operating the physical archive, and using ECS custom software for monitoring archive functions and maintaining the stored data. The Archive Manager may also work with the Automated Cartridge Storage Library System (ACSL) software and the AMASS Graphical User Interface (GUI). Finally, the Archive Manager conducts archive troubleshooting and problem resolution procedures.

Subsequent sections related to Archive Processing address procedures for the following functions:

- Section 17.1 Starting and Stopping AMASS.
- Section 17.2 Loading, removing, and managing archive media.
- Section 17.3 Monitoring and managing the archive with ECS custom GUIs.
- Section 17.4 Deleting granules from the archive.
- Section 17.5 Backing up and restoring AMASS.
- Section 17.6 Backing up and restoring archived data.
- Section 17.7 Archive troubleshooting.
- Section 17.8 ACSLS procedures.
- Section 17.9 Using the AMASS GUI.

For each set of functions, an **Activity Checklist** table provides an overview of the tasks to be completed. The outline of the Activity Checklist is as follows:

Column one - **Order** shows the order in which tasks could be accomplished.

Column two - **Role** lists the Role/Manager/Operator responsible for performing the task.

Column three - **Task** provides a brief explanation of the task.

Column four - **Section** provides the Procedure (P) section number or Instruction (I) section number where details for performing the task can be found.

Column five - **Complete?** is used as a checklist to keep track of which task steps have been completed.

17.1 Starting and Stopping AMASS

To start AMASS, the Archive Manager or System Administrator first ensures that the physical storage system is powered up and then enters commands at the FSMS server host (e.g., e0drg11, g0drg01, l0drg01, n0drg01) to start AMASS. Stopping AMASS is accomplished by killing the required daemons. Rebooting AMASS involves killing the daemons and then restarting the application.

Table 17.1-1 provides an Activity Checklist for Starting and Stopping AMASS.

Table 17.1-1. Starting and Stopping AMASS - Activity Checklist

Order	Role	Task	Section	Complete?
1	System Administrator or Archive Manager	Starting the AMASS Application	(P) 17.1.1	
2	System Administrator or Archive Manager	Shutting Down AMASS Tape Archive System	(P) 17.1.2	
3	System Administrator or Archive Manager	Rebooting AMASS	(P) 17.1.3	

17.1.1 Starting the AMASS Application

Starting the AMASS FSMS requires actions to ensure that the STK Powderhorn storage system is powered up as well as actions at the SGI FSMS host. Powering up the STK requires actions at its control panels, including the Library Management Unit (LMU) and Library Control Unit (LCU) [the Library Storage Module (LSM) is powered through the LCU]. *Note:* Preconditions include that 1) the FDDI network is up and running and 2) power to all units is functional and available.

Table 17.1-2 presents the steps required to start the AMASS application. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the

system, or have not performed this task recently, you should use the following detailed procedure:

- 1** Make sure power switches for the StorageTek LCU and LMU are **ON**.
 - *NOTE:* The LCU should be the last unit powered up, but otherwise there are no dependencies within the group.
- 2** If it is not already running, boot the FSMS SGI host (workstation **x0drg##**) normally.
 - *NOTE:* The **x** in the workstation name will be a letter designating your site: **g** = GSFC, **m** = SMC, **l** = LaRC, **e** = EDC, **n** = NSIDC, **o** = ORNL, **a** = ASF, **j** = JPL; the **##** will be an identifying two-digit number (e.g., **n0drg01** indicates an FSMS SGI host at NSIDC).
 - There are no dependencies on other hosts, COTS or custom software.
 - AMASS normally starts automatically on bootup. If it does, go to **Step 5**. If it does not, or if you are restarting AMASS after a shutdown, go to **Step 3**.
- 3** At the FSMS SGI host, log in as **root**.
- 4** Type **/usr/amass/tools/amass_start** and then press the **Return/Enter** key.
 - The AMASS application starts.
- 5** To verify that AMASS has started correctly, type **/usr/amass/bin/amassstat -c** and then press the **Return/Enter** key.
 - The message **FILESYSTEM IS ACTIVE** is displayed.

Table 17.1-2. Starting AMASS

Step	What to Do	Action to Take
1	Power switches ON .	Observe/set switches
2	Boot FSMS SGI host	Normal workstation boot
3	Log in as root	press Return/Enter
4	amass_start	press Return/Enter
5	amassstat -c	press Return/Enter

17.1.2 Shutting Down AMASS Tape Archive System

Table 17.1-3 presents the steps required to shut down AMASS. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in as **root** (system administrator) at the FSMS SGI host (workstation **x0drg##**).
 - NOTE: The **x** in the workstation name will be a letter designating your site: **g** = GSFC, **m** = SMC, **l** = LaRC, **e** = EDC, **n** = NSIDC, **o** = ORNL, **a** = ASF, **j** = JPL; the **##** will be an identifying two-digit number (e.g., **n0drg01** indicates an FSMS SGI host at NSIDC).
- 2 Type **/usr/amass/tools/killdaemons** and then press the **Return/Enter** key.
 - A message is displayed indicating that all daemons have been terminated.

Table 17.1-3. Shutting Down AMASS

Step	What to Do	Action to Take
1	Log in as root	Press Return/Enter
2	killdaemons	Press Return/Enter

17.1.3 Rebooting AMASS

The AMASS file system may need to be rebooted during certain anomalous conditions (e.g., system "hang," interruption of communication between AMASS and ACSLS, a required daemon is down). AMASS needs to have the following daemons running at all times: amassmain, daemons/lm_ip -a fslock, klogd, amass_iocomp, qset, libsched, libio_tape,. To verify they are running, simply search for the AMASS processes (refer to Procedure 17.7.1.1 **Checking Daemons and Using healthcheck**). To check the health of AMASS while it is still running, execute the **healthcheck** command (refer to Procedure 17.7.1.1).

In order to reboot AMASS you must have root privileges. The following procedure demonstrates the steps to reboot AMASS. Table 17.1-4 presents the steps required to follow the reboot process. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in as **root** (system administrator) at the FSMS SGI host (workstation **x0drg##**).
 - NOTE: The **x** in the workstation name will be a letter designating your site: **g** = GSFC, **m** = SMC, **l** = LaRC, **e** = EDC, **n** = NSIDC, **o** = ORNL, **a** = ASF, **j** = JPL; the **##** will be an identifying two-digit number (e.g., **n0drg01** indicates an FSMS SGI server at NSIDC).

- 2 To kill the daemons, type **killdaemons** and then press the **Return/Enter** key.
 - A message is displayed indicating that all daemons have been terminated.
- 3 If you want to test AMASS before restarting, go to step 4; otherwise, type **amass_start** and then press the **Return/Enter** key.
 - The AMASS application starts.
- 4 To test the AMASS filesystem prior to starting AMASS type: **install_tests**, and press the **Return/Enter** key.
 - Tests the operation jukebox operation and cache partitions, then restarts AMASS.

Table 17.1-4. Rebooting AMASS

Step	What to Do	Action to Take
1	Log in as root	press Return/Enter
2	killdaemons	press Return/Enter
3	amass_start	press Return/Enter
4	install_tests	press Return/Enter

17.2 Loading, Removing, and Managing Archive Media

For the STK storage facility, each Powderhorn is equipped with a 21-tape Cartridge Access Port (CAP). In automatic mode, tapes may be placed in the CAP for automatic loading. Tapes are also ejected through the CAP when identified for ejection using a command at the host for the STK Automated Cartridge System Library Software (ACSLs). It is also possible to bypass the CAP and manually load media directly into the library bins, typically only done at the initial load of the system or if it is otherwise necessary to load large numbers of volumes. Newly loaded volumes may need to be placed online and formatted. It is also necessary to ensure the ready availability of drive cleaning cartridges in the specially designated volume group for that purpose.

Table 17.2-1 provides an Activity Checklist for Loading, Removing, and Managing Archive Media.

Table 17.2-1. Loading, Removing, and Managing Archive Media - Activity Checklist

Order	Role	Task	Section	Complete?
1	Archive Manager	Automatically Loading Archive Media	(P) 17.2.1	
2	Archive Manager	Manually Loading Archive Media	(P) 17.2.2	
3	Archive Manager	Formatting a Volume	(P) 17.2.3	
4	Archive Manager	Removing Archive Media	(P) 17.2.4	

17.2.1 Automatically Loading Archive Media

Automatic loading of media is appropriate when there are relatively small numbers of media to be loaded. Up to 21 volumes at a time may be loaded through the Cartridge Access Port (CAP). With automated loading, AMASS assigns each cartridge a unique volume number, enters the volumes in its database, and marks the volumes Online in the database.

Table 17.1-2 presents the steps required for automated media loading. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in as **amass** or **root** at the FSMS SGI host (workstation **x0drg##**).
 - NOTE: The **x** in the workstation name will be a letter designating your site: **g** = GSFC, **m** = SMC, **l** = LaRC, **e** = EDC, **n** = NSIDC, **o** = ORNL, **a** = ASF, **j** = JPL; the **##** will be an identifying two-digit number (e.g., **n0drg01** indicates an FSMS SGI server at NSIDC).
- 2 At the FSMS host, type **/usr/amass/bin/bulkinlet SP** and then press the **Return/Enter** key.
 - The Cartridge Access Port (CAP) door unlocks (audible unlatching sound).
 - **Note:** If you have removed an existing volume and are re-inserting it, do not use the **SP** option, which puts the volume in the general space pool. Instead type **/usr/amass/bin/bulkinlet <volgrp>**, where **<volgrp>** is the volume group from which the volume was removed. This will put the volume back where it was before removal.
- 3 Write down or note the bar code number(s) on the label(s) of the cartridge(s), open the recessed latch on the CAP door and insert the tape(s), solid black side up, with the bar code label facing you, and close the door.
 - The robot scans all the volumes.
 - Data for the newly inserted media are displayed, including bar codes, associated volume numbers, and, in the **flag** column, the letters **IUO**, indicating that the volumes are inactive (**I**), unformatted (**U**), and offline (**O**).

- 4 For any newly inserted media, it is necessary to issue a formatting command. For the new 9940 tapes, type **/usr/amass/bin/volformat -b 256k ###**, where **###** is the volume number, and then press the **Return/Enter** key. You can enter more than one, separating each number from the preceding one with a space.
 - A message requests confirmation that you wish to continue.
- 5 Type **y** and then press the **Return/Enter** key.
 - A message is displayed requesting further confirmation, stating that **The following volumes will be formatted:** and listing volume numbers, followed by **(Y-N)**.
- 6 Type **y** and then press the **Return/Enter** key.
 - After a few minutes, a message **Completed formatting all volumes** is displayed.
- 7 To verify that the volume(s) are inserted, type **/usr/amass/bin/vollist** and then press the **Return/Enter** key.
 - Data for the media are displayed; the **flag** column shows that the newly formatted volumes are inactive (**I**).
- 8 To activate the media for use, type **/usr/amass/bin/volstat** and then press the **Return/Enter** key.
 - Data for the media are displayed; the **flag** column shows that the volumes are now active (**A**).

Table 17.2- 2. Automatically Loading Archive Media

Step	What to Do	Action to Take
1	Log in as amass or root	enter text; press Return/Enter
2	bulkinlet SP (unless re-inserting removed volume)	press Return/Enter
3	Place cartridge(s) in CAP	close door
4	volformat -b 256k < volumenumber >	press Return/Enter
5	y (to continue)	press Return/Enter
6	y (to confirm/continue)	press Return/Enter
7	vollist	press Return/Enter
8	volstat	press Return/Enter

17.2.2 Manually Loading Archive Media

Media may be introduced into volume groups in the storage facility without AMASS initial monitoring and assignment. This may be done using the CAP, as illustrated in the following procedure, or it may be done during an initial loading of the system. For such an initial loading,

large numbers of cartridges may be placed directly in storage slots without using the CAP (i.e., with the Powderhorn library door open before the system is powered up). Manual loading uses an AMASS command different from that used for automatic loading; the command used here enables AMASS to determine what media have been placed in the library and convey the information to its database.

Table 17.2-3 presents the steps required for manual media loading. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 To manually insert a tape into the Powderhorn, login to the control software (ACSL) using the **acssa** account at an ACSL workstation (e.g., e0drs03, g0drs03, l0drs02, n0drs03).
- 2 Type **enter 0,0,0** and then press the **Return/Enter** key.
 - The Cartridge Access Port (CAP) door unlocks (audible unlatching sound).
- 3 Write down or note the bar code number(s) on the label(s) of the cartridge(s), open the recessed latch on the Cartridge Access Port (CAP) door and insert the tape(s), solid black side up, with the bar code label facing you, and close the door.
 - The robot scans all the volumes.
- 4 At the AMASS host, type **/usr/amass/bin/bulkload -s SP** and then press the **Return/Enter** key.
 - The AMASS database is populated with data for the volumes in the library, including bar codes, associated volume numbers, and status -- inactive (**I**), unformatted (**U**), and offline (**O**). The data may be reviewed using the **vollist** command.
 - *Note:* If you are loading a very large number of volumes, such as at initial load, and choose to bypass the CAP and place the volumes directly in the library slots, data about the volumes will not be immediately available to ACSL for communication to AMASS. You will first have to use the ACSL **audit** command to initiate an audit of the library, a process that may take several hours.

Caution

Inactivate AMASS before using the following command.

- 5 To view a list of media in the library, type **/usr/amass/utills/medialist -3**, and then press the **Return/Enter** key.
 - The **-3** option indicates the STK Powderhorn.
 - The utility reads the library element status stored in the library, and information about the library contents, including the status (**FULL** or **EMPTY**) of the elements.

Table 17.2-3. Manually Loading Archive Media

Step	What to Do	Action to Take
1	Log in as acssa	enter text; press Return/Enter
2	enter 0,0,0	press Return/Enter
3	Place cartridge(s) in CAP	close door
4	bulkload -s SP	press Return/Enter
5	medialist -3	press Return/Enter

17.2.3 Formatting a Volume

To format a volume, it must be online. A volume is placed online using the **volloc** command. For a tape cartridge, you must first set the tape length using the **tapelength** command. Formatting a volume destroys any files on that volume. Before formatting a volume, check to make sure it does not have any files that should be saved. Table 17.2-4 presents the steps required to follow the formatting process. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 To put the volume online, at the FSMS host, type **/usr/amass/bin/volloc -n ###**, and then press the **Return/Enter** key.
 - **###** is the number of the volume.
- 2 To verify there are no files on volume, type **/usr/amass/bin/volfilelist <Vol. No.>** , and then press the **Return/Enter** key.
 - No files are displayed.
 - If a list of files is returned, indicating that the volume is not empty, before proceeding verify that you have the correct volume and that it is to be formatted.
- 3 To format the volume, type **/usr/amass/bin/volformat -b 256k ###**, and then press the **Return/Enter** key.
 - **###** is the number of the volume.
- 4 To verify status of the volume, type **/usr/amass/bin/volprint -a #####**, and then press the **Return/Enter** key.
 - **###** is the number of the volume.

Table 17.2-4. Formatting a Tape Volume

Step	What to Do	Action to Take
1	volloc -n volnumber	press Return/Enter
2	volfilelist volnumber	press Return/Enter
3	volformat -b 256k volnumber	press Return/Enter
4	volprint -a volnumber	press Return/Enter

17.2.4 Removing Media

Table 17.2-5 presents the steps required to remove media from the STK Powderhorn. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in as **amass** or **root** at the FSMS SGI host (workstation **x0drg##**).
 - NOTE: The **x** in the workstation name will be a letter designating your site: **g** = GSFC, **m** = SMC, **l** = LaRC, **e** = EDC, **n** = NSIDC, **o** = ORNL, **a** = ASF, **j** = JPL; the **##** will be an identifying two-digit number (e.g., **n0drg01** indicates an FSMS SGI server at NSIDC).
- 2 Determine which volumes you want to remove by utilizing the volume number. If necessary to review volume numbers and other information, at the FSMS host, type **/usr/amass/bin/vollist** and then press the **Return/Enter** key.
 - A list of volumes is displayed.
- 3 If there are only a few volumes to remove, for each volume to be removed type **/usr/amass/bin/voloutlet ###**, where **###** is the volume number, and then press the **Return/Enter** key.
 - AMASS marks the volume off-line and the volume is transferred to the CAP.
- 4 Open the recessed latch on the Cartridge Access Port (CAP) door and remove the tape(s).

Table 17.2-5. Removing Media from the Storage Library

Step	What to Do	Action to Take
1	Log in as amass or root	enter text; press Return/Enter
2	vollist	press Return/Enter
3	voloutlet volumenumber	press Return/Enter
4	Remove tape(s) from CAP	open CAP latch

17.3 Monitoring and Managing the Archive with ECS Custom GUIs

Custom Graphical User Interfaces (GUIs) in the ECS software can provide helpful information concerning the relationship between physical storage archives (Library Storage Modules, or LSMs) and the Archive Server software applications at the site. For example, a data repository identified as DRP1 is served by the software application EcDsStArchiveServerDRP1.

Subdivisions within LSMs (e.g., for storage of different data types) are reflected in the Storage Management database, where each Volume Group (a logical group of volumes in the archive) has its own path. Each path maps to an AMASS volume group, and thus to a physical volume group in the archive.

Information concerning archive servers and the logical volume groups served may be obtained from the Storage Management Control GUI. The Storage Configuration tab on the Storage Management GUI permits display of server information and access to related status information.

Table 17.3-1 provides an Activity Checklist for Monitoring and Managing the Archive with ECS Custom GUIs.

Table 17.3-1. Monitoring and Managing the Archive - Activity Checklist

Order	Role	Task	Section	Complete?
1	Archive Manager	Launching DSS GUIs	(P) 17.3.1	
2	Archive Manager	Using Storage Management GUIs to Display Archive Path Information	(P) 17.3.2	
3	Archive Manager	Monitoring Archive Requests using the Storage Management GUI	(P) 17.3.3	
4	Archive Manager	Monitoring Distribution Requests using the Data Distribution GUI	(P) 17.3.4	
5	Archive Manager	Setting Checksum Calculation	(P) 17.3.5	

17.3.1 Launching DSS GUIs

The following software applications are associated with DSS:

- Science Data Server (SDSRV).
- Storage Management (STMGT) Servers.
 - Request Manager Server.
 - Staging Disk Server.
 - Cache Manager Server.
 - Archive Server.
 - Request Manager Server.
 - FTP Server.

- D3/9940 Tape Server.
- 8mm Tape Stacker Server.
- Data Distribution (DDIST) Server.
- DDIST Graphical User Interface (GUI).
- STMGT GUIs.
- Science Data Server GUIs.

Access to Storage Management, Data Distribution (DDIST), and other GUIs is gained through the use of UNIX commands. The procedure for launching the GUIs begins with the assumption that the applicable servers are running and that the operator (Archive Manager or System Administrator) has logged in.

Table 17.3-2 presents the steps required to launch DSS GUIs using UNIX commands. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

1 Access the command shell.

- The command shell prompt is displayed.

NOTE: Commands in Steps 2 through 9 are typed at a UNIX system prompt.

2 Type **setenv DISPLAY *clientname*:0.0** and then press the **Return/Enter** key.

- Use either the terminal/workstation IP address or the machine-name for the *clientname*.

3 Start the log-in to the DDIST client server by typing **/tools/bin/ssh *hostname*** (e.g., **e0dis02**, **g0dis02**, **l0dis02**, or **n0dis02**) and then press the **Return/Enter** key.

- If you receive the message, **Host key not found from the list of known hosts. Are you sure you want to continue connecting (yes/no)?** type **yes** (“y” alone does not work).
- If you have previously set up a secure shell passphrase and executed **sshremote**, a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears; continue with Step 4.
- If you have not previously set up a secure shell passphrase; go to Step 5.

4 If a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears, type your *Passphrase* and then press the **Return/Enter** key. Go to Step 6.

5 At the **<user@remotehost>'s password:** prompt, type your *Password* and then press the **Return/Enter** key.

- 6 To change directory to the directory containing the startup scripts for DSS, type **cd /usr/ecs/<MODE>/CUSTOM/utilities** and then press the **Return/Enter** key.
 - The <MODE> will most likely be one of the following operating modes:
 - OPS (for normal operation).
 - TS1 or TS2 (for testing).
 - Note that the separate subdirectories under /usr/ecs apply to different operating modes.
- 7 To launch the Storage Management Control GUI, type the following command: **EcDsStmgtGuiStart <MODE>**, where <MODE> is the one selected in Step 6, and then press the **Return/Enter** key.
 - The Storage Management Control GUI, used for review of storage events and status of devices, is displayed.
- 8 To launch the Data Distribution GUI, use a similar procedure and type the following command: **EcDsDdistGuiStart <MODE>**, where <MODE> is the one selected in Step 6, and then press the **Return/Enter** key.
 - The Data Distribution GUI is displayed.
- 9 To launch the **DSS Science Data Server** GUI, log in to the host for Science Data Server (e.g., **e0acs05**, **g0acs03**, **l0acs03**, or **n0acs04**). Use a similar procedure and type the following command: **EcDsSdSrvGuiStart <MODE>** and then press the **Return/Enter** key.
 - The Science Data Server Operator GUI is displayed.

Table 17.3-2. Launching DSS GUIs

Step	What to Do	Action to Take
1	Access command shell	
2	Log in to appropriate host (for DDIST or SDSRV)	Enter text, press Return/Enter
3	cd /usr/ecs/<MODE>/CUSTOM/utilities	press Return/Enter
4	Go to Step 5 for STMGT, 6 for DDIST, or 7 for SDSRV	
5	EcDsStmgtGuiStart <MODE>	press Return/Enter
6	EcDsDdistGuiStart <MODE>	press Return/Enter
7	EcDsSdSrvGuiStart <MODE>	press Return/Enter

17.3.2 Using Storage Management GUIs to Display Archive Path Information

If requested to provide archive path information for a particular Earth Science Data Type (ESDT) stored in the archive, the Storage Management GUI can be used to obtain the needed

information. Table 17.3-3 presents the steps required to use the Storage Management GUIs to Display Archive Path information. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DSS Storage Management GUI using UNIX commands (see Procedure 17.3.1 **Launching DSS GUIs**).
 - The DSS Storage Management GUI is displayed.
- 2 Click on the **Storage Config.** tab to ensure that the Storage Configuration display is available.
 - The **Storage Config.** tab is displayed.
- 3 In the field listing **Server Type**, click on the **ARCHIVE** line to highlight it.
 - The selected line is highlighted and the **Server Name** and **Status** of archive servers are displayed in the field listing **Server Name**.
- 4 Click on the **Vol Grp Config.** tab.
 - The **Volume Group Information** is displayed showing volume groups and their current paths.
- 5 If it is desirable to display the path history for a data type, on the **Vol Grp Config.** tab, click on the **Data Type Name** entry for the specific data type for which path history information is desired.
 - The selected line is highlighted.
- 6 Click on the **Display History** button.
 - A **Volume Group History** window is displayed showing the path history for the highlighted data type.

Table 17.3-3. Using Storage Management GUIs to Display Archive Path Information and History

Step	What to Do	Action to Take
1	Launch the DSS Storage Management GUI	Use procedure in Section 17.3.1
2	Select Storage Config. tab	single-click
3	Highlight Archive in Server Type field	single-click
4	Select Vol Grp Config. tab	single-click
5	Highlight a selected data type line	single-click
6	Activate Display History button	single-click

17.3.3 Monitoring Archive Requests using the Storage Management GUI

A primary GUI tool for monitoring of archive processing is the **Request Status** window, accessible from the **Storage Management Control** GUI. Using the **Request Status** tab the Archive Manager or Distribution Technician can detect stalled requests or servers that appear to be idle.

The **Request Status** window displays the following information:

- **Operation** is the type of operation represented by the request.
- **Request ID** is a unique identifier for the request.
- **Progress** is the stage of processing on which the request is currently working (may include a numeric progress indication).
- **Status** provides information about processes attempted and the outcome (e.g., DsEstDRExecuteFailed, DsEstARPathSearchExhausted, OK, . . .WriteFailed, . . .).
- **Priority** is **Xpress**, **Very High**, **High**, **Normal**, or **Low**.
- **When Submitted** is the time and date when the request was received by the Storage Management server that is responsible for the request.
- **Last Updated** is the time and date when the status was last updated for the request.

The operator can reduce the displayed list of requests by clicking on the **Filtering** pull-down menu just above the **Request Status Information** list on the window. This permits filtering on four areas or filter types selectable from the pull-down menu:

- **Server** controls what activity is displayed by limiting the list to the requests being/having been serviced by a specific server. Selecting **All** displays all requests throughout Storage Management. Other selections include the individual archive servers, cache manager servers, ftp servers, request manager server, and staging disk servers.
- **Operation** allows the operator to focus on a specific type of operation. The list of operations is dynamically generated to reflect those operations for which requests are currently in queue (e.g., **All**, **CMLink**, **ArStore**, **ArRetrieve**, **FtpPull**, **FtpPush**).
- **Processing State** allows the operator to differentiate among requests that are being actively processed; have been completed, either successfully or to a retryable error state; or have been suspended and are awaiting the outcome of another event. The following selections are available: **All**, **Processing**, **Suspended**, **Completed**.
- **Submitter** allows the Distribution Technician to see the status of requests submitted by a specific client process. The list of possible clients is dynamically generated to reflect the list of clients with outstanding requests (e.g., **All**, **DSDD**, **HDFC**, **SDSV**, **this**, [various servers]).

Table 17.3-4 presents the steps required to monitor archive requests using the Storage Management Control GUI. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DSS Storage Management GUI using UNIX commands (see Procedure 17.3.1 **Launching DSS GUIs**).
 - The DSS Storage Management GUI is displayed.
- 2 Click on the **Request Status** tab.
 - The **Request Status** tab is displayed.
- 3 Observe information displayed on the **Request Status** tab of the **Storage Management Control** GUI.
 - The **Request Status Information** table displays the following information:
 - Operation.
 - Request ID.
 - Progress.
 - Status.
 - Priority.
 - When Submitted.
 - Last Updated.
 - By default all storage management server requests for the last 24 hours are shown in the **Request Status Information** table of the **Request Status** tab.
 - Clicking on any of the column headers of the **Request Status Information** table causes the listed requests to be sorted in order by the column selected.
 - For example, clicking on the **Last Updated** column header causes requests to be listed in order from the least recently updated to the most recently updated.
 - The **Operator Messages** field at the bottom of the GUI displays messages concerning events occurring in storage management operations.
 - Note that storage management servers control virtually all data inserted into or retrieved from the archive; the resulting large amount of activity on the **Request Status** tab may make it useful to restrict the number of requests displayed by applying a filter (see next step).

- 4 To filter the list of requests, use the **Filtering** pull-down menu above the top left corner of the **Request Status Information** table, selecting as desired to display requests associated with a particular **Server**, **Operation**, **Processing State**, or **Submitter**.
 - The list of requests displayed in the **Request Status Information** table is restricted by the filtering choice.
- 5 Observe the Storage Management requests displayed in the **Request Status Information** table.
 - The **Progress** and **Status** column entries in the table may provide indication for particular requests of potential problems or conditions requiring attention.
- 6 Repeat Steps 4 and 5 as necessary to monitor Storage Management requests.
- 7 To **exit**, follow menu path **File→Exit**.

Table 17.3-4. Monitoring Archive Requests using the Storage Management GUI

Step	What to Do	Action to Take
1	Launch the DSS Storage Management GUI	Use procedure in Section 17.3.1
2	Select Request Status tab	single-click
3	Observe listed request information	read text
4	Select Filtering option	click-hold and drag cursor or execute three clicks
5	Observe selected (filtered) requests	read text
6	Repeat steps 4 and 5 as necessary	
7	Exit (if necessary/desirable)	Menu selection File→Exit

17.3.4 Monitoring Distribution Requests using the Data Distribution GUI

Distribution requests result from orders for ECS data, placed by users or subscriptions, and by requests for data by internal ECS processes (e.g., those related to data processing). As ECS responds to these requests, the Archive Manager or other operators can monitor the progress of the distribution requests. Table 17.3-5 presents the steps required to monitor distribution requests using the Data Distribution GUI. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the **DDIST GUI** using UNIX commands (see Procedure 17.3.1 **Launching DSS GUIs**).
 - The **Data Distribution GUI** tool is displayed.

- 2 Click on the **Distrib'n Requests** tab.
 - The Distribution Requests window is opened.
 - A list of requests is displayed.
- 3 Follow menu path **View→Filter**.
 - The **Distribution Filter Requests** window opens.
 - Three filter types are displayed in a radio box at the top of the **Distribution Filter Requests** window: **Request ID**, **Requester**, and **All Requests**.
- 4 In the radio box at the top, click on one of the radio buttons to select filtering by **Request ID**, **Requester**, or **All Requests** as desired.
 - The selection is indicated by the radio button depressed appearance.
 - If the selection is **Request ID** or **Requester**, the cursor moves to the text entry field to the right of the selected button.
- 5 If the selection is **Request ID** or **Requester**, enter the request ID or requester's name, respectively, in the appropriate text entry field.
- 6 In the **Media Type:** area of the **Distribution Filter Requests** window, click on the **All** button or click on one of the entries in the **Media Type:** field to select for filtering on **FtpPull** or **FtpPush** (because of the incorporation of the Product Distribution System to handle media distributions, any media distribution requests are reflected as FtpPush distributions to the Product Distribution System).
 - Selected entries in the **Media Type:** window show as highlighted.
- 7 In the **State:** area, click on the **All** button to select all states, or click on one or more radio buttons to select one or more states for the filtering.
 - Any selected **State:** toggle buttons show as depressed.
- 8 Click on the **OK** push button, located at the bottom of the window.
 - The other push buttons located at the bottom of the window are Apply, Cancel, and Help.
 - The Filter Requests window is closed.
 - The Distribution Requests screen shows any requests that meet the filter criteria in the **Data Distribution Requests** field.
- 9 If necessary, use the scroll bar at the bottom of the **Data Distribution Requests** field to scroll horizontally to view the state of the selected request(s).

Table 17.3-5. Monitoring Distribution Requests using the Data Distribution GUI

Step	What to Do	Action to Take
1	Launch the DSS Data Distribution GUI	Use procedure in Section 17.3.1
2	Select Distrb'n Requests tab	single-click
3	Open the Distribution Filter Requests window	Menu selection View→Filter
4	Select Filtering type	single-click
5	Enter any required information for Filtering type	enter text
6	Select Media Type: filtering option(s)	click to highlight
7	Select State: filtering option(s)	click to select
8	Click O K to apply filters and close Distribution Filtering Requests window	single-click
9	Observe filtered Distribution Request(s) list, scrolling as necessary	read text; click-hold and drag as necessary

17.3.5 Setting Checksum Calculation

The system design for ECS incorporates calculation of a checksum when a granule is inserted into the archive. If such a checksum is calculated, it can then be used as an indicator to determine if there is data corruption within the archive. Comparison of the original checksum with one calculated, for example, when the granule is retrieved (e.g., for processing or distribution) can detect whether the inserted file and the retrieved file are the same. If the checksums do not match, then the operator can investigate (refer to Procedure 17.7.4 **Diagnosing/Investigating Read Errors**). The checksums are set in the configuration for the archive server, with variables that set calculation on granule insert and calculation on retrieval. The Storage Management GUI provides an easy way to set these configuration parameters. The settings are available from the **Storage Config.** tab, by highlighting the Archive Server and clicking on the **Modify Server** button. This opens the **Archive Server Configuration** window. The window includes option buttons to **Enable Checksumming On Store:** and **Enable Checksumming On Retrieve:**

Calculation of checksums can be time consuming. System throughput may be significantly improved if checksum calculation on granule insert is turned off, and therefore the default reflects checksum calculation turned off. Unfortunately, turning checksums off compromises the ability to detect data corruption in the archive. This problem may be alleviated somewhat by calculating a checksum when a granule is first retrieved from the archive and storing that checksum to be compared with one calculated upon a later retrieval. However, this approach will not guard against the possibility of data corruption on initial insertion (e.g., through I/O errors). If it becomes necessary to enable calculation (e.g., for troubleshooting), use the Storage Management GUI. Table 17.3-6 presents the general steps required for setting checksum calculation. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the Storage Management GUI using UNIX commands (refer to Procedure 17.3.1 **Launching DSS GUIs**).
 - The DSS Storage Management GUI is displayed.

- 2 Click on the **Storage Config.** tab to ensure that the Storage Configuration display is available.
 - The **Storage Config.** tab is displayed.
- 3 In the field listing **Server Type**, click on the **ARCHIVE** line to highlight it.
 - The selected line is highlighted and the **Server Name** and **Status** of archive servers are displayed in the field listing **Server Name**.
- 4 In the field listing **Server Name**, click on the archive server (e.g., **EcDsStArchiveServerDRP1**) for which the checksum variables are to be set.
 - The selected line is highlighted.
- 5 Click on the **Modify Server** button.
 - The **Archive Server Configuration** window is displayed showing configuration data for the selected archive server.
- 6 Click on the option button in the **Enable Checksumming On Store:** block.
 - A pop-up display offers the choice of **Yes** or **No**.
- 7 Click on the desired choice to enable (**Yes**) or disable (**No**) checksumming when a granule is stored.
 - The pop-up display is closed and the selected choice appears as the label on the option button.
- 8 Click on the option button in the **Enable Checksumming On Retrieve:** block.
 - A pop-up display offers the choice of **Yes** or **No**.
- 9 Click on the desired choice to enable (**Yes**) or disable (**No**) checksumming when a granule is retrieved.
 - The pop-up display is closed and the selected choice appears as the label on the option button.
- 10 Click on the **OK** button.
 - The **Archive Server Configuration** window is closed and the changes take effect.

Table 17.3-6. Setting Checksum Calculation

Step	What to Do	Action to Take
1	Launch the DSS Storage Management GUI	Use procedure in Section 17.3.1
2	Select Storage Config. tab	single-click
3	Select ARCHIVE server type	single-click to highlight
4	Select archive server for which to set checksumming	single-click to highlight
5	Select Modify Server button	single-click
6	Display options for Enable Checksumming on Store:	single-click
7	Select Yes to enable or No to disable	single-click
8	Display options for Enable Checksumming on Retrieve:	single-click
9	Select Yes to enable or No to disable	single-click
10	Select OK button	single-click

17.4 Deleting Granules

The **Granule Deletion** capability allows operators to delete products on demand. There are several circumstances that may require deletion on demand, such as:

- New PGE versions have been created and are used to reprocess large amounts of past data, creating new ESDT versions. As reprocessing progresses, Operations deletes the granules for the old ESDT versions from the archive and inventory.
- It is determined that certain lower-level (e.g., Level 2) products are of little or no interest to the science or public user community. In concert with the science teams, DAAC operations personnel decide to remove these products from the inventory. Since the products are still referenced by higher-level products as inputs, the DAAC decides to keep the inventory records for production history purposes.
- One or more granules were found defective and were reprocessed on an individual basis. When the reprocessing is complete, the operator wishes to delete the old, defective granule(s) from the inventory.
- A DAAC has extended ECS with subsetting services. The subsetting products are produced outside ECS, but are then inserted into the ECS archive to take advantage of the ECS distribution capability. The DAAC writes a script to delete the subsetting products on a regular basis.

The Science Data Server provides an application programming interface (API) for deleting granules from the archive, or from both the archive and inventory since earlier releases, but the Granule Deletion capability adds a front-end command-line utility that provides several ways for selecting granules for deletion. Confirmation is generally required so that granules are not inadvertently deleted. However, the confirmation may be suppressed so that operators can run regularly scheduled deletion scripts using background execution. This suppression possibility presents an opportunity for inadvertent loss of data and so must be used with care and only after thorough testing of any deletion script.

The Science Data Server captures deletions and related errors in the application log. Operators may also specify a separate and independent delete log for immediate analysis of the success or failure of a delete operation.

Deletion Sequence. The deletion of granules from the archive involves three elements, and therefore actually occurs in stages. Two of the elements are parts of the Science Data Server (SDSRV), and the third is a part of the Storage Management (STMGT) software and Graphical User Interface (GUI).

- *Logical Deletion:* For the first stage, a command-line delete utility specifies selection criteria for deletion of granules and "logically" deletes from the inventory those granules that meet the criteria. These granules are flagged as 'deleted' and can no longer be accessed, but their inventory entries are not yet removed. The logical 'deletion' may specify, via command line input, removal of granule files from the archive (*Delete From Archive*, or DFA) only, leaving the inventory record, or it may specify *Physical Deletion*, which entails removal of the inventory record as well as removal of the files from the archive. The deletion flag consists of records in the SDSRV database. Specifically, in the DsMdGranules table, the value of the DeleteFromArchive entry is changed from **N** to **Y**, and the granule is entered in the DsMdDeletedGranules table with a time stamp recording the logical deletion time.
- *Physical Deletion:* The second stage is actual deletion from the inventory of those granules marked for physical deletion (not DFA only), which occurs when the operations staff runs the physical deletion cleanup utility script. For Physical Deletion, the script removes all inventory rows for granules that were flagged as 'deleted,' including rows referencing related information (e.g., QA data). The script writes to the STMGT database (and therefore must be run under a log in by *sdsrv_role* with authorization to write to that database), creating entries in the DsSdPendingDelete table for granules to be deleted. This includes entries for granules that are to be physically deleted, as well as those designated DFA only. The operations staff controls the lag time between logical deletion and physical deletion. That lag time is entered into the physical deletion script, which deletes only inventory entries for granules that have been logically deleted prior to that time period.
- *Deletion from Archive (DFA):* STMGT provides a GUI screen that allows the operator to initiate the removal from the archive of the files listed its deletion table (populated by SDSRV). STMGT creates requests to the archive servers to delete files. The STMGT GUI can be used to look at the state of the deletion requests. Files that are successfully deleted have their associated rows removed from the STMGT database table.

Periodically, as sufficient data removal from the archive makes it appropriate, operations may elect to reclaim the tape space and recycle archive tapes. The AMASS software commands (*volcomp*, *volclean*, *volformat*, *volstat*) are used for that purpose.

Table 17.4-1 provides an Activity Checklist for Deleting Granules from the Archive.

Table 17.4-1. Deleting Granules from the Archive - Activity Checklist

Order	Role	Task	Section	Complete?
1	Archive Manager/Database Administrator	Resetting the Lock on the DsMdDeletedGranules table	(P) 17.4.1	
2	Archive Manager	Selecting Granules for Deletion	(P) 17.4.2	
3	Archive Manager/ Data Base Administrator	Deleting Granules from the Inventory and Archive (Physical Deletion)	(P) 17.4.3	
4	Archive Manager	Deleting Granules from the Archive	(P) 17.4.4	

17.4.1 Resetting the Lock on the DsMdDeletedGranules Table

The DsMdDeletedGranules table may become locked during execution of the Deletion Cleanup task if there is a need to restart the Science Data Server or if there is a problem with Sybase. The lock can prevent granules being marked for deletion upon subsequent runs of the Granule Deletion utility, and it is therefore necessary to reset the lock. This is accomplished using a Perl script, **EcDsResetLock.pl**. Table 17.4-2 presents the steps required to run the EcDsResetLock.pl script. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in to the SDSRV host (e.g., e0acs05, g0acs03, l0acs03, n0acs04) with an ID authorized with permissions to execute the EcDsResetLock.pl script.
- 2 To change directory to the directory containing the script, type **cd /usr/ecs/<MODE>/CUSTOM/utilities** and then press the **Return/Enter** key.
 - The **<MODE>** will most likely be one of the following operating modes:
 - OPS (for normal operation).
 - TS1 or TS2 (for testing).
 - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/utilities**.
- 3 To execute the script, type **EcDsResetLock.pl** and then press the **Return/Enter** key.
 - The script prompts **Enter Mode of Operation :**.
- 4 To identify the mode in which the script is to be run, type **OPS**, **TS1**, **TS2**, or other **<MODE>**, and then press the **Return/Enter** key.
 - The script prompts **Enter Log File name:**.
- 5 Type **ResetLock.log** or another name of your choosing for the log, and then press the **Return/Enter** key.
 - The script prompts **Enter Sybase User Name:**.

- 6 Type **sdsrv_role** and then press the **Return/Enter** key.
 - The script prompts **Enter Sybase password:**.
- 7 Type the Sybase **<password>** (may require input from Database Administrator) and then press the **Return/Enter** key.
 - The script prompts **Enter Sybase SQL Server Name:**.
- 8 Type **x0acs0n_srvr** and then press the **Return/Enter** key.
 - The **x** will be **g** for GSFC, **e** for EDC, **l** for LaRC, or **n** for NSIDC. The **n** will be a number identifying the Sybase SQL Server (may require input from the Database Administrator).
 - The script prompts **Enter SDSRV's database name:**.
- 9 Type **EcDsScienceDataServer1_<MODE>**, and then press the **Return/Enter** key.
 - The script attempts to access the locked table, generating a number of errors, before resetting the lock so that marking for deletion can proceed.

Table 17.4-2. Resetting the Lock on the DsMdDeletedGranules Table

Step	What to Do	Action to Take
1	Log in to SDSRV host	enter text
2	cd /usr/ecs/<MODE>/CUSTOM/utilities	enter text; press Return/Enter
3	EcDsResetLock.pl	enter text; press Return/Enter
4	<MODE>	enter text; press Return/Enter
5	ResetLock.log	enter text; press Return/Enter
6	sdsrv_role	enter text; press Return/Enter
7	<password>	enter text; press Return/Enter
8	x0acs0n_srvr	enter text; press Return/Enter
9	EcDsScienceDataServer1_<MODE>	enter text; press Return/Enter

17.4.2 Selecting Granules for Deletion

Selecting granules for deletion consists of running the Granule Deletion Command line tool to accomplish *logical deletion*, or marking (by entries in the SDSRV inventory database) of granules for deletion. It specifies either DFA only, leaving the inventory record, or Physical Deletion, which will result in removal of the inventory record as well as removal of the files from the archive.

17.4.2.1 Selection by ESDT ShortName, Version, and Granule Time Coverage

Table 17.4-3 presents the steps required to select granules for deletion using the ESDT ShortName, version, and granule time coverage. If you are already familiar with the procedure,

you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 10 Log in to the SDSRV host (e.g., e0acs05, g0acs03, l0acs03, n0acs04) with an ID authorized with permissions to execute the Granule Deletion utility.
- 11 To change directory to the directory containing the script, type **cd /usr/ecs/<MODE>/CUSTOM/utilities** and then press the **Return/Enter** key.
 - The **<MODE>** will most likely be one of the following operating modes:
 - OPS (for normal operation).
 - TS1 or TS2 (for testing).
 - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/utilities**.
- 12 Type the command to run the Granule Deletion utility specifying ESDT ShortName, ESDT version, and granule time coverage (granule beginning date and time, and granule ending date and time), and specifying DFA or physical deletion, and then press the **Return/Enter** key.
 - For deletion from archive only, the command is:
EcDsGranuleDeleteClientStart <MODE> -name <ShortName> -version <version no.> -BeginDate <granbegdate/time> -EndDate <granenddate/time> -log /usr/ecs/<MODE>/CUSTOM/logs/GranDel<n>.log -DFA.
 - For physical deletion, the command is:
EcDsGranuleDeleteClientStart <MODE> -name <ShortName> -version <version no.> -BeginDate <granbegdate/time> -EndDate <granenddate/time> -log /usr/ecs/<MODE>/CUSTOM/logs/GranDel<n>.log -physical.
 - The utility executes and displays the number of granules for deletion, and prompts the user **Do you want to continue [y/n]?**.
 - **Note:** It is possible to suppress the continuation prompt by including the argument **-noprompt** with the command.

13 Type **y** and then press the **Return/Enter** key.

- The process continues to completion.
- **Note:** The deletion actions are displayed in the Deletion log and in the Science Data Serve ALOG, including information on the user ID of the requester, the ShortName, VersionID, and granule coverage time of the request. It is also possible to view the SDSRV database to verify the granule tagging for deletion; the granule should appear in the database with values depending on whether the deletion request specified **-DFA** or **-physical** (see matrix below).

Request Type	DsMdGranules Table		DsMdDeletedGranules Table	
	DeleteEffective Date	DeleteFromArchive	Transaction Time	DFA Flag
-DFA	NULL	Y	Current Time	1
-physical	Current Time	N	Current Time	0

Table 17.4-3. Selection by ESDT ShortName, Version, and Granule Time Coverage

Step	What to Do	Action to Take
1	Log in to SDSRV host	enter text
2	cd /usr/ecs/<MODE>/CUSTOM/utilities	enter text; press Return/Enter
3	Run Granule Delete utility	Execute Step 3 of Procedure 17.4.2.1
4	y (at continuation prompt y/n?)	enter text; press Return/Enter

17.4.2.2 Selection by ESDT ShortName, Version, and Granule Insert Time Range

Table 17.4-4 presents the steps required to select granules for deletion using the ESDT ShortName, version, and granule insert time range. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 14 Log in to the SDSRV host (e.g., e0acs05, g0acs03, l0acs03, n0acs04) with an ID authorized with permissions to execute the Granule Deletion utility.
- 15 To change directory to the directory containing the script, type **cd /usr/ecs/<MODE>/CUSTOM/utilities** and then press the **Return/Enter** key.
- The **<MODE>** will most likely be one of the following operating modes:
 - OPS (for normal operation).
 - TS1 or TS2 (for testing).
 - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/utilities**.

- 16 Type the command to run the Granule Deletion utility specifying ESDT ShortName, ESDT version, and granule insert time range (insert beginning date and time, and insert ending date and time), and specifying DFA or physical deletion, and then press the **Return/Enter** key.

- For deletion from archive only, the command is:

EcDsGranuleDeleteClientStart <MODE> -name <ShortName> -version <version no.> -insertbegin <insbegdate/time> -insertend <insenddate/time> -log /usr/ecs/<MODE>/CUSTOM/logs/GranDel<n>.log -DFA.

- For physical deletion, the command is:

EcDsGranuleDeleteClientStart <MODE> -name <ShortName> -version <version no.> -insertbegin <insbegdate/time> -insertend <insenddate/time> -log /usr/ecs/<MODE>/CUSTOM/logs/GranDel<n>.log -physical.

- The utility executes and displays the number of granules for deletion, and prompts the user **Do you want to continue [y/n]?**.
- **Note:** It is possible to suppress the continuation prompt by including the argument **-noprompt** with the command.

- 17 Type **y** and then press the **Return/Enter** key.

- The process continues to completion.
- **Note:** The deletion actions are displayed in the Deletion log and in the Science Data Serve ALOG, including information on the user ID of the requester, the ShortName, VersionID, and granule insert time of the request. It is also possible to view the SDSRV database to verify the granule tagging for deletion; the granule should appear in the database with values depending on whether the deletion request specified **-DFA** or **-physical** (see matrix below).

Request Type	DsMdGranules Table		DsMdDeletedGranules Table	
	DeleteEffective Date	DeleteFromArchive	Transaction Time	DFA Flag
-DFA	NULL	Y	Current Time	1
-physical	Current Time	N	Current Time	0

Table 17.4-4. Selection by ESDT ShortName, Version, and Insert Time Range

Step	What to Do	Action to Take
1	Log in to SDSRV host	enter text
2	cd /usr/ecs/<MODE>/CUSTOM/utilities	enter text; press Return/Enter
3	Run Granule Delete utility	Execute Step 3 of Procedure 17.4.2.2
4	y (at continuation prompt y/n?)	enter text; press Return/Enter

17.4.2.3 Selection Using a Separate Input File

The tool permits referencing a list of granules in a file created for the purpose of providing that list as input to the tool. The desired deletion is achieved by creating a file containing identifying information for the granule(s) to be deleted. The file can list multiple granules. Several options are available for identifying granules, and for controlling how the utility executes:

- the input file can specify SDSRV Granule IDs (geoIDs).
- the input file can specify logical, or "local," Granule IDs.
- the command to execute the Granule Deletion utility can include the **-display** argument, which results in a listing of the geoID and logical ID of each granule selected for deletion.
- by default, the number of granules selected for deletion is displayed and the operator is prompted to confirm the deletion, but the command to execute the Granule Deletion utility can include a **-noprompt** argument, which suppresses the confirmation prompt.
- the operator can choose whether selected granules are to be deleted from the archive and the inventory or from the archive only.
- by default, any BROWSE, QA, and PH granules associated with physically deleted granules are deleted if no longer referenced otherwise, but the command to execute the Granule Deletion utility can include a **-noassoc** argument, which suppress deletion of these associated granules.

Table 17.4-5 presents the steps required to select granules for deletion using a separate input file. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 18 Log in to the SDSRV host (e.g., e0acs05, g0acs03, l0acs03, n0acs04) with an ID authorized with permissions to execute the Granule Deletion utility.
- 19 To change directory to the directory containing the script, type **cd /usr/ecs/<MODE>/CUSTOM/utilities** and then press the **Return/Enter** key.
- The **<MODE>** will most likely be one of the following operating modes:
 - OPS (for normal operation).
 - TS1 or TS2 (for testing).
 - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/utilities**.
- 20 Type the command to run the Granule Deletion utility specifying an input file type and file name, a granule deletion log file path, type of deletion (DFA or physical), omission of prompt (if prompt is not desired), argument to suppress deletion of associated BROWSE, QA, and PH granules (if deletion of those associated granules is not desired), and then press the **Return/Enter** key.
- The form of the command is:
- ```
EcDsGranuleDeleteClientStart <MODE> [-geoidfile <geoidfilename>]
[-localgranulefile <localfilename>] -log /usr/ecs/<MODE>/CUSTOM/logs/
GranDel<n>.log -DFA [-noprompt] [-noassoc]
```
- OR
- ```
EcDsGranuleDeleteClientStart <MODE> [-geoidfile <geoidfilename>]
[-localgranulefile <localfilename>] -log /usr/ecs/<MODE>/CUSTOM/logs/
GranDel<n>.log -physical [-noprompt] [-noassoc].
```
- The utility executes and displays the number of granules for deletion, and, unless the command included the **-noprompt** option, prompts the user **Do you want to continue [y/n]?**. If the **-noprompt** option was used, the process continues to completion (see *Note* under step 4).
- 21 If prompted at the end of Step 3, type **y** and then press the **Return/Enter** key.
- The process continues to completion.
 - **Note:** The deletion actions are displayed in the Deletion log and in the Science Data Serve ALOG, including information on the user ID of the requester, the ShortName, VersionID, and granule insert time of the request. It is also possible to view the SDSRV database to verify the granule tagging for deletion; the granules should appear in the database with values depending on whether the deletion request specified **-DFA** or **-physical** (see matrix below). By default, if there are any BROWSE, QA, and PH granules associated with the science data granules to be deleted and these associated granules are not referenced by other granules, the

associated granules will also be tagged for deletion. However, the command may include the **-noassoc** argument to suppress deletion of these associated granules.

Request Type	DsMdGranules Table		DsMdDeletedGranules Table	
	DeleteEffective Date	DeleteFromArchive	Transaction Time	DFA Flag
-DFA	NULL	Y	Current Time	1
-physical	Current Time	N	Current Time	0

Table 17.4-5. Selection Using a Separate Input File

Step	What to Do	Action to Take
1	Log in to SDSRV host	enter text
2	cd /usr/ecs/<MODE>/CUSTOM/utilities	enter text; press Return/Enter
3	Run Granule Delete utility	Execute Step 3 of Procedure 17.4.2..3
4	If prompted, y (at continuation prompt y/n?)	enter text; press Return/Enter

17.4.3 Deleting Granules from the Inventory and Archive (Physical Deletion)

Once granules are marked for deletion, the operator runs the Deletion Cleanup Utility, a Perl script, **EcDsDeletionCleanup.pl**. This script identifies those granules that were marked for deletion to the STMGT database for removal from the archive. In addition, if the granules were marked for physical deletion in the SDSRV database, they are deleted from the SDSRV database. A lag time, specified as a number of days, is used so that the operator can request that not all the granules marked for deletion are immediately deleted -- i.e., the Deletion Cleanup Utility deletes granules marked for deletion a specified number of days prior to the current date. A lag time of 0 may be used to implement immediate deletion. The script also asks for entry of a batch size, by which the operator specifies the increments for transfer of files from the SDSRV database to the STMGT database. For large numbers of deletions, a large batch size may be specified (up to a maximum of 10,000).

Table 17.4-6 presents the steps required to delete granules from the inventory and archive. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 22** Log in to the SDSRV host (e.g., e0acs05, g0acs03, l0acs03, n0acs04) with an ID authorized with permissions to execute the Deletion Cleanup utility.
- 23** To change directory to the directory containing the script, type **cd /usr/ecs/<MODE>/CUSTOM/utilities** and then press the **Return/Enter** key.
 - The **<MODE>** will most likely be one of the following operating modes:
 - OPS (for normal operation).
 - TS1 or TS2 (for testing).

- The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/utilities**.
- 24 To run the Deletion Cleanup Utility, type **EcDsDeletionCleanup.pl** and then press the **Return/Enter** key.
- The script prompts **Enter lag time in days:**.
- 25 Type **<n>**, where **n** is the lag time in days specifying a number of days prior to the current date during which granules marked for deletion are not yet to be deleted, and then press the **Return/Enter** key.
- **Note:** Use a lag time of 0 to indicate immediate deletion of all marked granules.
 - The script prompts **Enter mode of operation:**.
- 26 Type **<MODE>**, where **<MODE>** is the mode in which you are making the deletion (typically **OPS**, **TS1**, or **TS2**) and then press the **Return/Enter** key.
- The script prompts **Enter log file name:**.
- 27 Type **DelCleanup1.log** and then press the **Return/Enter** key.
- The script prompts **Enter Sybase User:**.
- 28 Type **sdsrv_role**, and then press the **Return/Enter** key.
- The script prompts **Enter Sybase User Password:**.
- 29 Type **<password>**, where **<password>** is the Sybase password (**Note:** This step may require input by the Database Administrator).
- The script prompts **Enter sql server:**.
- 30 Type **<x>0acg<nn>_srvr** (**e0acg11_srvr** at EDC, **g0acg01_srvr** at GSFC, **l0acg02_srvr** at LaRC, or **n0acg01_srvr** at NSIDC) and then press the **Return/Enter** key.
- The script prompts **Enter DBName:**.
- 31 Type **EcDsScienceDataServer1_<MODE>**, where **<MODE>** is the mode in which you are making the deletion (typically **OPS**, **TS1**, or **TS2**) and then press the **Return/Enter** key.
- The script prompts **Enter STMGT DBName:**.
- 32 Type **stmgtddb1_<MODE>**, where **<MODE>** is the mode in which you are making the deletion (typically **OPS**, **TS1**, or **TS2**) and then press the **Return/Enter** key.
- **Note:** If you do not know the STMGT DBName, contact the Database Administrator.
 - The script prompts **Enter Batch size:**.

- 33 Type **<n>** where **n** is the batch size specifying the increments for transfer of files from the SDSRV database to the STMGT database (maximum 10,000), and then press the **Return/Enter** key.
- The script prompts **Do you wish to continue deleting these granules?:**.
- 34 Type **y** and then press the **Return/Enter** key.
- The Deletion Cleanup Utility script displays the number of granules to be deleted from the archive (DFA) and physically deleted, with a confirmation prompt. If the lag time was specified as 0, all granules in the DeletedGranules table are displayed.
- 35 Type **y** (response to the confirmation prompt) and then press the **Return/Enter** key.
- Execution of the Deletion Cleanup Utility script completes.
 - **Note:** In the SDSRV database, the SDSRV Staging table (DsMdStagingTable) can be observed for transfer of data to the STMGT database (in increments of the specified batch size); when the transfer is complete, the table is empty. In the STMGT database, the STMGT Pending Delete table (DsStPendingDelete) can be observed for receipt of the data; all granules specified in the delete request are received. The Deletion Cleanup log displays messages about the actions, indicating that information is placed in the STMGT database in increments of the specified batch size.

Table 17.4-6. Deleting Granules from the Inventory and Archive

Step	What to Do	Action to Take
1	Log in to SDSRV host	enter text
2	cd /usr/ecs/<MODE>/CUSTOM/utilities	enter text; press Return/Enter
3	EcDsDeletionCleanup.pl	enter text; press Return/Enter
4	<n> (lag time, in days)	enter text; press Return/Enter
5	<MODE>	enter text; press Return/Enter
6	DelCleanup.log	enter text; press Return/Enter
7	sdsrv_role	enter text; press Return/Enter
8	<password>	enter text; press Return/Enter
9	x0acg0nn_srvr	enter text; press Return/Enter
10	EcDsScienceDataServer1_<MODE>	enter text; press Return/Enter
11	stmgtdb1_<MODE>	enter text; press Return/Enter
12	<n> (increments for transfer of files, maximum 10,000)	enter text; press Return/Enter
13	y (to continue)	enter text; press Return/Enter
14	y (to confirm/continue)	enter text; press Return/Enter

17.4.4 Deleting Granules from the Archive

Once the STMGT database receives the data on the granules to be deleted (reflected in the STMGT Pending Delete table, DsStPendingDelete), the operator uses the Storage Management GUI to initiate the removal from the archive of the listed files. This completes the physical

deletion for those granules selected for removal from the archive and inventory, and accomplishes the removal from the archive of those granules selected for Deletion from the Archive (DFA) only. Table 17.4-7 presents the steps required to delete granules from the archive. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DSS Storage Management GUI using UNIX commands (see Procedure 17.3.1 **Launching DSS GUIs**).
 - The DSS Storage Management GUI is displayed.
- 2 On the STMGT GUI, to view the ESDTs with granules targeted for deletion, follow menu path **Delete→Batch Delete**.
 - The **Batch Delete** window is displayed, listing the granules tagged for deletion as ESDT/Version pairs with numbers of files in the **Granule Deletion Information** field.
- 3 To select data for deletion from the archive, click on an ESDT/Version pair.
 - The selected ESDT/Version pair is highlighted.
- 4 Click the **Delete** button.
 - A confirmation Delete Warning prompt asks **Are you sure you want to delete the selected files?**
- 5 To confirm the deletion, click the **OK** button.
 - The delete request continues to completion and the ESDT/Version pair is removed from the list of granules tagged for deletion in the **Granule Deletion Information** field of the **Batch Delete** window.
 - **Note:** The delete actions can be tracked via messages in the Archive Server log files (EcDsStArchiveServer.ALOG, EcDsStArchiveServerDebug.log)

Table 17.4-7. Deleting Granules from the Archive

Step	What to Do	Action to Take
1	Launch the DSS Storage Management GUI	Use procedure in Section 17.3.1
2	Open the Batch Delete window	Menu selection Delete→Batch Delete
3	Select an ESDT/Version pair	single-click
4	Activate Delete button	single-click
5	Click OK (to confirm deletion)	single-click

17.5 Backing Up and Restoring AMASS

A key responsibility of the Archive Manager is to guard against loss of the AMASS database and functioning. This is achieved through creation of backups that can be used to restore functioning in the event of database corruption or other failure. The archive storage format used by AMASS is a proprietary format designed to optimize storage and retrieval speed. The command **vgexport -q** can be used to create a text file, storable on magnetic media, which can be used with the AMASS format archive tapes and the command **vgimport** to recover from the loss. This command exports the AMASS database for a specified volume group to standard out (**stdout**), a file containing the directory structure and media attributes (e.g., media type, ownership, timestamp) for the volume group. The file is located in **/usr/amass/filesysdb** and is exported as standard ASCII text.

Table 17.5-1 provides an Activity Checklist for activities related to backing up and restoring AMASS.

Table 17.5-1. Activity Checklist for Backing Up and Restoring AMASS

Order	Role	Task	Section	Complete?
1	Archive Manager/ System Administrator	Creating a Backup for AMASS	(P) 17.5.1	
2	Archive Manager/ System Administrator	Replacing the AMASS Database Backup Volume (Volume 1)	(P) 17.5.2	
3	Archive Manager/ System Administrator	Restoring the AMASS Database	(P) 17.5.3	

17.5.1 Creating a Backup for AMASS

Table 17.5-2 presents the steps required to create a backup for AMASS. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01) as **amass** or **root**.
- 2 Type **/usr/amass/bin/vgexport -q** and then press the **Return/Enter** key.
 - A file named **stdout** is created in **/usr/amass/filesysdb**.
 - *Note:* The **stdout** file is useful only with the archive volumes represented in the AMASS database.

Table 17.5-2. Creating a Backup for AMASS

Step	What to Do	Action to Take
1	Log in as amass or root	enter text; press Return/Enter
2	vgexport -q	enter text; press Return/Enter

17.5.2 Replacing the AMASS Database Backup Volume (Volume 1)

The AMASS database backup is stored in the archive on Volume 1. "Volume 1," hard coded to be the backup volume, actually designates one of the last volumes in the StorageTek Library Storage Module, to prevent its inadvertent use as a data volume. Whenever **amassbackup** is run, AMASS issues an e-mail message with information on volume capacity and usage. It is also possible to issue the command **vollist 1** to display how much space is left on the volume, or **volprint 1** for still more detail. If the volume becomes full *during* a backup attempt, the backup will fail and it is necessary to initialize a new backup volume and perform a full backup as described in the following procedure. Table 17.5-3 presents the steps required to replace the AMASS database backup volume. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01) as **amass** or **root**.
- 2 Type **/usr/amass/bin/voloutlet 1** and then press the **Return/Enter** key.
 - The LSM robot places the Backup Volume in the CAP.
- 3 Open the recessed latch on the CAP door; remove the Backup Volume tape and store it in a safe place.
- 4 Physically designate the new Backup Volume tape so that it can be easily discriminated from other volumes (e.g., write "Backup Volume" on the tape, color code the tape, or make and display a note of its home storage slot or preprinted barcode).
- 5 Note the pre-printed number on the volume label (e.g., 112102), insert the new Backup Volume in the CAP, and close the door.
 - The robot scans the volume.
- 6 At the AMASS host, type **/usr/amass/bin/bulkinlet -u** and then press the **Return/Enter** key.
 - AMASS assigns the Backup Volume a unique volume number.
 - AMASS marks the volume **ONLINE** in the AMASS database.
 - AMASS assigns the Backup Volume to the last barcode position in the library.
 - AMASS gives the volume a **BACKUP VOLUME** label.

7 Type **/usr/amass/bin/vollist 1**, and then press the **Return/Enter** key.

- AMASS displays the following:

VOL NUM	VOL GRP	JUKE NUM	POS	VOL LABEL	FLAGS	USED (MB)	AVAIL (MB)	DEAD (%)	ERRS
1	0	1		BACKUP-VOLUME	I	0	20000	0	0

8 To change the Volume Label field from **BACKUP-VOLUME** to the preprinted media number (e.g., 112102), type **/usr/amass/bin/vollabel 1 112102** and then press the **Return/Enter** key.

9 Type **/usr/amass/bin/vollist 1**, and then press the **Return/Enter** key.

- AMASS displays the following:

VOL NUM	VOL GRP	JUKE NUM	POS	VOL LABEL	FLAGS	USED (MB)	AVAIL (MB)	DEAD (%)	ERRS
1	0	1		112102	I	0	20000	0	0

10 Type **/usr/amass/bin/volformat -u** and then press the **Return/Enter** key.

- A message requests confirmation that you wish to continue.

11 Type **y** and then press the **Return/Enter** key.

- A message is displayed requesting further confirmation, stating that **The following volumes will be formatted: 1 (Y-N)**.

12 Type **y** and then press the **Return/Enter** key.

- After a few minutes, a message **Completed formatting all volumes** is displayed.

13 To verify that the volume is inserted, type **/usr/amass/bin/vollist 1** and then press the **Return/Enter** key.

- Data for the media are displayed; the **flag** column shows that the newly formatted volume is inactive (**I**).

14 Type **/usr/amass/bin/amassbackup -fv** and then press the **Return/Enter** key.

- AMASS performs a full backup with the verbose option of the AMASS database and transaction logs.

Table 17.5-3. Replacing the AMASS Database Backup Volume

Step	What to Do	Action to Take
1	Log in as amass or root	enter text; press Return/Enter
2	voloutlet 1	enter text; press Return/Enter
3	Remove Backup Volume	open CAP door
4	Physically designate Backup Volume	mark volume
5	Note pre-printed number on volume label	read label
6	bulkinlet -u	enter text; press Return/Enter
7	vollist 1	enter text; press Return/Enter
8	vollabel 1 nnnnnn (number from step 5)	enter text; press Return/Enter
9	vollist 1	enter text; press Return/Enter
10	volformat -u	enter text; press Return/Enter
11	y (to continue)	enter text; press Return/Enter
12	y (to confirm volume to be formatted)	enter text; press Return/Enter
13	vollist 1	enter text; press Return/Enter
14	amassbackup -fv	enter text; press Return/Enter

17.5.3 Restoring the AMASS Database

The AMASS database is restored manually by the System Administrator or the Archive Manager using the AMASS command **amassrestore**. This command restores the last full backup, the last partial backup, and all journal transactions that have occurred since the last backup. It creates a sub-directory under filesysdb called **journal**. All restored files are copied to the **journal** directory. The following restore procedure uses a backup volume or tape device. Table 17.5-4 presents the steps required to restore the AMASS database. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01) as **amass** or **root**.

Caution

Do not use the **amassrestore** command when AMASS is running. To shutdown AMASS, refer to the Special Shutdown Procedures in the AMASS technical documentation *Installing AMASS*.

- 2 To inactivate the AMASS file system, type **/usr/amass/bin/amassstat -i**.
 - The AMASS file system is inactivated.
- 3 Make sure the backup drive is available.
 - If there is another volume in the drive, return it to its home slot by entering **/usr/amass/daemons/amassrecovery -s** (the option **-s** prevents system startup and performs file recovery).

4 Type `/usr/amass/bin/amassrestore -v -L <barcodelabel>` and then press the **Return/Enter** key.

- If you do not know the barcode label number for the backup volume, it can be obtained by entering `/usr/amass/bin/vollist 1`.
- The AMASS database is restored from the backup volume.

Table 17.5-4. Restoring the AMASS Database

Step	What to Do	Action to Take
1	Log in as amass or root	enter text; press Return/Enter
2	amassstat -i	enter text; press Return/Enter
3	Ensure backup drive is available (if necessary, <code>/usr/amass/daemons/amassrecovery -s</code>)	enter text; press Return/Enter
4	amassrestore -v -L <barcodelabel>	enter text; press Return/Enter

17.6 Backing Up and Restoring Archived Data

The ECS archive design incorporates programmed backups of archived data. System requirements specify that a percentage of archived data be duplicated for local and offsite storage to provide for data safety. However, the large volume of ECS archived data merits finding alternatives to complete backup of all volumes in the libraries. Selection of data for backup is based on assessment of the feasibility of recovery in the event of data loss.

It is imperative to back up data that would be irretrievable if lost. Such data are saved to the archive, saved to local backup, and saved to offsite backup. Many data elements that will be archived, however, could be retrieved in the event of loss. For example, in the event of loss of a higher level product that is an output of processing a lower level product, it would be possible to restore the higher level product by reprocessing the lower level product. As another example, ECS will often archive a lower level product from a data provider, but that product may also be retained in the archives of the data provider. If the product were lost from the ECS archive, it would be possible to ingest it again from the data provider, using appropriate Ingest procedures.

Thus, when data are inserted into the archive (e.g., through Ingest, from Processing), up to three copies of the data may be created, reflecting different types of data use:

- the active archive copy, available for distribution or other use (volume group is specified in the *Archive ID*).
- a copy to be retained for local backup (volume group is specified in the *Backup ID*).
- a copy to be sent to offsite backup storage (volume group is specified in the *Offsite ID*).

The Archive Manager has the responsibility for ensuring and managing necessary backups of archived data and, in the event of loss, executing or supporting efforts to recover lost data. Table 17.6-1 provides an Activity Checklist for backing up and restoring archived data.

Table 17.6-1. Activity Checklist for Backup and Restoration of Archived Data

Order	Role	Task	Section	Complete?
1	Archive Manager/ Science Data Specialist	Creating Offsite Backups	(P) 17.6.1.1	
2	Archive Manager/ System Administrator	Creating Replacement Backups Manually from Existing Archives	(P) 17.6.1.2	
3	Archive Manager	Manual Data Recovery from Local Backup Tapes	(P) 17.6.2.1	
4	Archive Manager	Manual Data Recovery from Offsite Backup Tapes	(P) 17.6.2.2	
5	Archive Manager	Manual Data Recovery from Damaged Cartridge	(P) 17.6.2.3	
6	Archive Manager	Data Recovery for Known Files	(P) 17.6.2.4	
7	Archive Manager	SDSRV Retrieval of File Location Metadata	(P) 17.6.2.5.1	
8	Archive Manager	SDSRV Retrieval of Granule Production History Metadata	(P) 17.6.2.5.2	

17.6.1 Backups for Archive Data

The paths for creation of the data copies are specified for each ESDT when it is loaded. The Archive ID (for the archive copy) and the Backup ID (for the local backup copy) should reflect different archives if possible (i.e., different Library Storage Modules), to spread the risk of loss. The Offsite ID will not be a remote site path, but rather a local path for making copies to be sent for offsite storage. The requirements to implement creation of offsite backups include:

- creating a subdirectory and volume group for offsite backups.
- using the **Vol Grp Config.** tab of the Storage Management GUIs to add the volume group to the appropriate archive server and set the offsite ID to be the three-character specification for the local site (e.g., EDC, GSF, LAR, NSC).
- adding volumes to the volume group as needed.

17.6.1.1 Creating Offsite Backups

Each site is responsible for arranging its own secure offsite storage. The offsite backup cartridges are removed from the archive storage facility using procedures already described (see Procedure 17.2.4). For local and/or offsite storage of specific archive data, the Archive Manager generates or directs the generation of a list of selected data. At the time the files are archived, they are written to specific and separate volume groups that correspond to the three data usage types identified previously (i.e., active archive, local backup, offsite backup). Only files belonging to the data usage type are written to the tapes in a specific volume group. The Archive Manager or Science Data Specialist sets up these volume groups when an Earth Science Data Type (ESDT) is installed. Table 17.6-2 presents the steps required to create offsite backups. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you

are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DSS Storage Management GUI using UNIX commands (see Procedure 17.3.1 **Launching DSS GUIs**).
 - The DSS Storage Management GUI is displayed.
- 2 Click on the **Vol Grp Config.** tab to display the Volume Group information.
 - The **Vol Grp Config.** tab information is displayed.
- 3 Click on the **Add . . .** button below the Volume Group Information field.
 - The **Add Volume Group** window is displayed.
- 4 In the **Add Volume Group** window, click in the **Data Type.Version:** field.
 - The cursor moves to the **Data Type.Version:** field.
- 5 Type the ESDT *ShortName* and *Version* (e.g., MOD01.001) of the data type for which the volume group is to be created.
 - The typed entry appears in the **Data Type.Version:** field.
- 6 In the **Add Volume Group** window, click on the pull-down arrow at the end of the **HWCI:** field.
 - A pull-down menu displays designators of the hardware configuration items available for storing data.
- 7 Click on the designator for the hardware configuration item where the archive copies of data for the ESDT are to be stored.
 - The selected designator is displayed in the **HWCI:** field.
- 8 In the **Add Volume Group** window, click in the **Volume Group Path:** field.
 - The cursor moves to the **Volume Group Path:** field.
- 9 Type the full path identification for the storage of active archive data for the ESDT (typically, the path will be of the form **dss_stkn/<MODE>/xxxxx**, where *n* is a number designating a StorageTek Library Storage Module, *MODE* is **OPS**, **TS1**, or **TS2**, and *xxxxx* is a short identifier for what is being stored; e.g., **dss_stk1/OPS/modl0**).
 - The typed entry appears in the **Volume Group Path:** field.
- 10 In the **Volume Group Type:** radio box, click on the **PRIMARY** button.
 - The button depressed appearance indicates selection of **PRIMARY**, signifying that the volume group being created is for primary storage for active archive use.

- 11 Click on the **Save and Add Next VG** button at the bottom of the **Add Volume Group** window.
 - The volume group is created for display in the **Volume Group Information** field on the **Vol Grp Config.** tab of the Storage Management GUI.
- 12 In the **Add Volume Group** window, click in the **Volume Group Path:** field.
 - The cursor moves to the **Volume Group Path:** field.
- 13 Change the data entered at Step 9 to identify the full path for the storage of local backup data for the ESDT.
 - *Note:* This step is only for those ESDTs that require local backup.
 - The typed entry appears in the **Volume Group Path:** field.
- 14 In the **Volume Group Type:** radio box, click on the **BACKUP** button.
 - The button depressed appearance indicates selection of **BACKUP**, signifying that the volume group being created is for storage for local backup use.
- 15 Click on the **Save and Add Next VG** button at the bottom of the **Add Volume Group** window.
 - The volume group is created for display in the **Volume Group Information** field on the **Vol Grp Config.** tab of the Storage Management GUI.
- 16 In the **Add Volume Group** window, click in the **Volume Group Path:** field.
 - The cursor moves to the **Volume Group Path:** field.
- 17 Change the data entered at Step 13 to identify the full path for the creation and initial storage of offsite backup data for the ESDT.
 - *Note:* This step is only for those ESDTs that require offsite backup.
 - The typed entry appears in the **Volume Group Path:** field.
- 18 In the **Volume Group Type:** radio box, click on the **OFFSITE** button.
 - The button depressed appearance indicates selection of **OFFSITE**, signifying that the volume group being created is for creation and initial storage for offsite backup use.
- 19 Click on the **Save and Exit** button at the bottom of the **Add Volume Group** window.
 - The volume group is created for display in the **Volume Group Information** field on the **Vol Grp Config.** tab of the Storage Management GUI.
 - Data stored in the volume group for OFFSITE backup can be safeguarded by removing tapes that have data stored on them (see procedure 17.2.4) and transporting the tapes to a secure offsite storage location.

Table 17.6-2. Creating Offsite Backups

Step	What to Do	Action to Take
1	Launch the DSS Storage Management GUI	Use procedure in Section 17.3.1
2	Select Vol Grp Config. tab	single-click
3	Click Add . . . button	single-click
4	Data Type.Version field	single-click
5	Enter ESDT ShortName and Version	enter text
6	Display HWCI: pull-down menu	single-click
7	Select HWCI:	single-click
8	Volume Group Path: field	single-click
9	Identify active archive path	enter text
10	Select PRIMARY volume group type	single-click
11	Activate Save and Add Next VG button	single-click
12	Volume Group Path: field	single-click
13	Identify local backup path	enter text
14	Select BACKUP volume group type	single-click
15	Activate Save and Add Next VG button	single-click
16	Volume Group Path: field	single-click
17	Identify offsite backup path	enter text
18	Select OFFSITE volume group type	single-click
19	Activate Save and Exit button	single-click

17.6.1.2 Creating Replacement Backups Manually from Existing Archives

If loss of data necessitates obtaining and inserting backup data from local or offsite storage, it is necessary to create replacement data to be returned to backup storage. Table 17.6-3 presents the steps required to create replacement backups manually from existing archives. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01) as **amass** or **root**.
- 2 Type **/usr/amass/bin/volcopy -c <source> <destination>** (where **<destination>** is the volume number of the destination volume and **<source>** is the volume number of the source volume), and then press the **Return/Enter** key.
 - The **-c** option specifies copy of the source to the destination.
 - A bit for bit copy of the source (the cartridge to be copied) is made at the destination (an available, unused cartridge). Because the copy procedure depends on the amount of data on the source cartridge, the process can take as long as an hour to complete.
 - *Note:* After starting a **volcopy** procedure, do not attempt to kill the process with the **kill -9** command.

- 3 A hardcopy/softcopy list of the files backed up should be created and kept for future file restoration operations.
- 4 Remove the backup volume(s) and send to offsite storage area, as appropriate.

Table 17.6-3. Creating Replacement Backups Manually from Existing Archives

Step	What to Do	Action to Take
1	Log in as amass or root	enter text; press Return/Enter
2	volcopy -c <source> <destination>	enter text; press Return/Enter
3	Add to list of backed up files for future reference	print list
4	Send backup to secure offsite storage	

17.6.2 Restoring Archive Data

Although the Archive hardware is highly reliable, errors due to tape or drive failure must be expected to occur, though at an extremely low rate, as a function of the archived data volume. Where errors have occurred and data has been lost from the archive and can not be restored from backup there may exist the potential to recover and re-archive equivalent data by one of the following means:

- copying from backup onto the original or a new primary.
- replacing damaged or corrupted volumes with vendor restored or backup volumes.
- re-generation by reprocessing.
- obtaining replacement data from the original external provider.

If a backup volume is available and contains the data that were lost or corrupted on the primary copy, the data can be copied using standard UNIX commands. If the backup volume must be obtained from offsite storage, it must then be inserted into the archive and brought on line. The procedures for loading archive media were addressed under a preceding topic. The requirements then entail:

- using the **Storage Config.** tab of the Storage Management GUIs to view the volume groups of the appropriate archive server and to find the files in the primary and backup volume groups.
- using the UNIX copy command (**cp** or **dd**) to copy the lost or corrupted file from the backup version to the primary version.
- as appropriate (i.e., if the recovery is one of a set of files to be restored, for example, because they were lost from a damaged tape), removing the names of the files recovered from the list of files to be recovered by other means.

If an entire volume is to be copied, perform the procedure to create replacement backups (see Procedure 17.6.1.2); if recovery is from offsite, send the backup back to secure offsite storage.

17.6.2.1 Manual Data Recovery from Local Backup Tapes

The procedure for manual data recovery from local backup tapes assumes that the tape is on-line and in the Powderhorn Library Storage Module. Volume groups and tapes are transparent to the automated file and storage management system. As long as the AMASS database is aware of the files, the operator moves data using standard UNIX commands.

Table 17.6-4 presents the steps required for manual data recovery from local backup tapes. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DSS Storage Management GUI using UNIX commands (see Procedure 17.3.1).
 - The DSS Storage Management GUI is displayed.
- 2 Click on the **Vol Grp Config.** tab to display volume group information.
 - The **Vol Grp Config.** tab is displayed.
- 3 Click in the **Find Next** field under the **Volume Group Information** field.
 - The cursor moves to the **Find Next** field.
- 4 Type the first three letters of the ESDT short name for the data type with missing or corrupted/damaged files on its primary storage tape.
 - The typed entry is displayed in the field.
- 5 Click on the **Find Next** button.
 - In the **Volume Group Information** field, the volume group information for the first volume group containing the three letters specified in Step 4 is highlighted.
- 6 As necessary, scroll further through the list of entries in the **Volume Group Information** field to locate the **Current Volume Group Path** for the primary and backup storage for the data type with missing or corrupted/damaged files on its primary storage tape; note or record the paths.
- 7 In a UNIX window, at the command line prompt, type **cp <backuppath/filename> <primarypath/filename>** and press the **Return/Enter** key.
 - The backup file is copied to the primary tape.

- 8 Repeat Step 7 as needed until all missing or corrupted/damaged files are restored from the backup tape to the primary tape.
- 9 Remove the file(s) restored in Steps 7 and 8 from any list of files to be recovered.

Table 17.6-4. Manual Data Recovery from Local Backup Tapes

Step	What to Do	Action to Take
1	Launch the DSS Storage Management GUI	Use procedure in Section 17.3.1
2	Select Vol Grp Config. tab	single-click
3	Find Next field	single-click
4	Specify search token	enter text
5	Activate Find Next button to start search	single-click
6	Locate and record primary and backup paths	click and drag to scroll; read text
7	cp <backuppath/filename> <primarypath/filename>	enter text; press Return/Enter
8	Repeat Step 7 as needed	
9	Remove restored files from list of files to be recovered	

17.6.2.2 Manual Data Recovery from Offsite Backup Tapes

Each site has its own arrangements for managing data requiring secure offsite backup storage. In the event of loss of data on primary and local backup tapes, recovery may be possible using offsite backup tapes. Table 17.6-5 presents the steps required for manual data recovery from local backup tapes. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DSS Storage Management GUI using UNIX commands (see Procedure 17.3.1).
 - The DSS Storage Management GUI is displayed.
- 2 Click on the **Vol Grp Config.** tab to display volume group information.
 - The **Vol Grp Config.** tab is displayed.
- 3 Click in the **Find Next** field under the **Volume Group Information** field.
 - The cursor moves to the **Find Next** field.
- 4 Type the first three letters of the ESDT short name for the data type with missing or corrupted/damaged files on its primary and local backup storage tapes.
 - The typed entry is displayed in the field.
- 5 Click on the **Find Next** button.
 - In the **Volume Group Information** field, the volume group information for the first volume group containing the three letters specified in Step 4 is highlighted.
- 6 As necessary, scroll further through the list of entries in the **Volume Group Information** field to locate the **Current Volume Group Path** for the primary and offsite storage for

the data type with missing or corrupted/damaged files on its primary and local backup storage tapes; note or record the paths.

7 Log in at the FSMS SGI host (workstation **x0drg##**).

- NOTE: The **x** in the workstation name will be a letter designating your site: **g** = GSFC, **m** = SMC, **l** = LaRC, **e** = EDC, **n** = NSIDC, **o** = ORNL, **a** = ASF, **j** = JPL; the **##** will be an identifying two-digit number (e.g., **n0drg01** indicates an FSMS SGI server at NSIDC).

8 To identify the offsite volume ID where a known file to be recovered is stored, on the FSMS host, at the command line prompt in a UNIX window, type **/usr/amass/utl/fileprint <filepathname/filename>**, where **filepathname/filename** is the path and name of the file, and press the **Return/Enter** key.

- AMASS returns database information for each location where the file is stored. For example, if the input filepathname and filename for a lost or damaged file is **/dss_stk1/aster/:Science:AST_L1BT:2137:1.EOSHDF**, the output returned by AMASS should look similar to the following:

```
FILE :Science:AST_L1BT:2137:1.EOSHDF :
      rid      =      5993
      prid     =      4749
      size     =      5410105 (0x528d39)
      start blk =      37750397
      vol      =       18
      ltvol    =       18
      mode     =      81a4
      links    =       1
      nrcrc    =      4195
      flags    =       0
```

This indicates that the file should be on volume 18. Similar output should be returned for each volume involved in storage of the file.

9 To determine if the offsite volume is in the archive, on the FSMS host, at the command line prompt in a UNIX window, type **/usr/amass/bin/vollist <volumenumber>**, where **volumenumber** is the volume ID returned in Step 8, and press the **Return/Enter** key.

- AMASS returns information about the requested volume. This step can be repeated for each volume ID returned in Step 8. If the return is similar to the following:

VOL NUM	VOL GRP	JUKE NUM	POS	VOL LABEL	FLAGS	USED (MB)	AVAIL (MB)	DEAD (%)	ERRS
18	700	1	NET	SD2102	IO	35589	0	35	0

the **IO** in the **FLAGS** column indicates that the volume is inactive and offline -- i.e., the volume is not in the Library Storage Module. The offsite backup volume will have this status if it is not in the archive and needs to be retrieved from offsite storage.

- 10 Retrieve the volume from offsite storage and insert it in the Library Storage Module (see Procedure 17.2.1), using the command **/usr/amass/bin/bulkinlet <volgrp>**, where *volgrp* is the identifier in the **VOL GRP** column of the return in Step 9.
 - The CAP door unlocks (audible unlatching sound).
- 11 Open the recessed latch on the CAP door and insert the tape(s), solid black side up, with the bar code label facing you, and close the door
 - The robot scans the volume(s) and makes the insertion into the volume group specified in Step 10.
- 12 To recover a file from the newly inserted offsite backup volume, in a UNIX window, at the command line prompt, type **cp <offsitepath/filename> <primarypath/filename>** and press the **Return/Enter** key.
 - The backup file is copied to the primary tape.
- 13 Repeat Step 12 as needed until all missing or corrupted/damaged files are restored from the offsite tape to the primary tape.
- 14 Remove the file(s) restored in Steps 12 and 13 from any list of files to be recovered.

Table 17.6-5. Manual Data Recovery from Offsite Backup Tapes

Step	What to Do	Action to Take
1	Launch the DSS Storage Management GUI	Use procedure in Section 17.3.1
2	Select Vol Grp Config. tab	single-click
3	Find Next field	single-click
4	Specify search token	enter text
5	Activate Find Next button to start search	single-click
6	Locate and record primary and offsite backup paths	click and drag to scroll; read text
7	Log in at FSMS host	enter text; press Return/Enter
8	/usr/amass/utls/fileprint <filepathname/filename>	enter text; press Return/Enter ; read text
9	/vollist <volumenumber>	enter text; press Return/Enter ; read text
10	bulkinlet <volgrp> volume retrieved from offsite	enter text; press Return/Enter
11	Insert the tape in the CAP	close door
12	cp <offsitepath/filename> <primarypath/filename>	enter text; press Return/Enter
13	Repeat Step 12 as needed	
14	Remove restored files from list of files to be recovered	

17.6.2.3 Manual Data Recovery from Damaged Cartridge

In the course of operations it is possible for a tape to become physically damaged or accidentally overwritten. Some indications of a damaged tape may be AMASS read/write errors, or AMASS may determine that the volume is unreadable and mark it inactive. In that event that a tape volume is damaged, a manual recovery of data from the cartridge must be attempted.

Because of the technical complexity of data recovery from a damaged cartridge, it will be performed by STK personnel. However, the Archive Manager can support and prepare for the process by listing all the files on the tape and their associated start block numbers and providing the list to the recovery personnel. The list is generated by using a *Perl* utility script. The utility will generate three ASCII files that must be provided to the STK recovery personnel along with the damaged tape. The files are: **filelist_<volnumber>**, **start_block_listing_volnumber**, and **README_<volnumber>**, where *volnumber* is the volume number of the requested tape volume.

The script utility, **EcDsStFilesPerTapeUtility**, is located in the directory **/usr/ecs/<MODE>/CUSTOM/utilities**. The script produces directory information followed by three files.

The directory information output should be similar to the following example:

```
/data1/data/:BR:Browse.001:1170:1.BINARY
/data1/data/:BR:Browse.000:1170:1.BINARY
/data1/data/:SC:MOD00:65001:1.CCSDS
/data1/data/:SC:MOD00:65002:1.CCSDS
/data1/data/:SC:MOD00:20001:1.CCSDS
/data1/data/:PH:PH.001:2000000076:1.BINARY
/data1/data/:PH:PH.000:2000000076:1.BINARY
/data1/data/:QA:QA.001:1003:1.ASCII
/data1/data/:QA:QA.001:1004:1.ASCII
/data1/data/:QA:QA.001:1005:1.ASCII
/data1/data/:OR:OR.001:2100:1.ASCII
/data1/data/:OR:OR.001:2101:1.ASCII
/data1/data/:OR:OR.001:2102:1.ASCII
/data1/data/:OR:OR.001:2103:1.ASCII
/data1/data/:AN:AN.001:3100:1.ASCII
/data1/data/:AN:AN.001:3101:1.ASCII
/data1/data/:AN:AN.001:3102:1.ASCII
/data1/data/:AN:AN.001:3103:1.ASCII
```

The information in the file **filelist_<volnumber>** is in ASCII format with one file name per line, as in the following example:

```
/dss_stk2/joel/TestStdSeq6_0_10.wrt
/dss_stk2/joel/TestStdSeq6_0_10.wrt
/dss_stk2/joel/TestStdSeq6_0_10.wrt
/dss_stk2/joel/TestStdSeq6_0_10.wrt
/dss_stk2/joel/TestStdSeq6_0_10.wrt
```

Table 17.6- 6 presents the steps required for manual data recovery from a damaged cartridge. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in as **amass** or **root** at the FSMS SGI host (workstation **x0drg##**, **xacg##**, or **xwkg##**).
 - NOTE: The **x** in the workstation name will be a letter designating your site: **g** = GSFC, **m** = SMC, **l** = LaRC, **e** = EDC, **n** = NSIDC, **o** = ORNL, **a** = ASF, **j** = JPL; the **##** will be an identifying two-digit number (e.g., **n0drg01** indicates an FSMS SGI server at NSIDC).
- 2 To verify that AMASS is running, type **/usr/amass/bin/amassstat -c** and press the **Return/Enter** key.
 - The message **FILESYSTEM IS ACTIVE** should be displayed. If it is not, restart AMASS using procedure 17.1.1.
- 3 To identify the volume ID where a known file to be recovered is stored, on the FSMS host, at the command line prompt in a UNIX window, type **/usr/amass/utls/fileprint <filepathname/filename>**, where *filepathname/filename* is the path and name of the file, and press the **Return/Enter** key.
 - AMASS returns database information for each location where the file is stored. For example, if the input filepathname and filename for a lost or damaged file is **/dss_stk1/aster/:Science:AST_L1BT:2137:1.EOSHDF**, the output returned by AMASS should look similar to the following:

```
FILE :Science:AST_L1BT:2137:1.EOSHDF :
      rid      =    5993
      prid      =    4749
      size      =   5410105 (0x528d39)
      start blk =   37750397
      vol       =     18
      ltvol     =     18
      mode      =    81a4
      links     =     1
      ncrs      =   4195
      flags     =     0
```
 - This indicates that the file should be on volume 18. Similar output should be returned for each volume involved in storage of the file.
- 4 Remove the volume from the archive (see Procedure 17.2.4).
- 5 Inspect the physical cartridge and tape for damage. Any creasing, scratches, snapping, or stretching of the tape may warrant keeping the volume offline and sending it to STK for replacement.

- 6 If the cartridge is damaged and to be returned to STK for recovery of data, run the script **EcDsStFilesPerTapeUtility** script; to start the utility script, type the command: **/usr/ecs/<MODE>/CUSTOM/utilities/EcDsStFilesPerTapeUtility** and press the **Return/Enter** key.
- The script runs and prompts for input of the volume number, as follows:
You have invoked an ECS utility script.

This script supports file recovery from an AMASS tape volume by generating two listings of the files located on that volume. The listings are ASCII files and can be viewed.

AMASS must be running in order to generate one of the listings. If a .fileprint. use error messages result, make sure AMASS is running, and you have AMASS privileges, before invoking this utility again.

Please enter the AMASS volume number,
for which you wish to generate listings

-->
- 7 Type the volume ID determined in Step 3 and press the **Return/Enter** key.
- The script runs and a message is displayed to indicate generation of the information and completion of the run. The two ASCII files are **filelist_<volnumber>** and **start_block_listing_volnumber>**. There is also a **README_<volnumber>** file.
- 8 Send the volume to STK along with the files generated by the perl utility.
- STK copies all uncorrupted data to a new tape and inserts filler data blocks to replace the lost data.
 - The filler data is inserted using the original block sequence so that the remaining data can be accessed by AMASS.
 - After copying of the data to a new cartridge, it is returned to the DAAC with the original volume label and a report indicating which data blocks were replaced with filler data.
- 9 After receiving the recovered tape back from STK, insert the tape into the library (see Procedure 17.2.1), using the command **/usr/amass/bin/bulkinlet <volgrp>**, where **volgrp** is the volume group number. (*Note:* If you do not know the volume group number, you can determine it by using the **vollist** command with the volume ID obtained in the return from Step 3).
- AMASS reads the volume label and places the volume in its home slot.

10 To put the volume **online**, type **volloc -n <volumenumber>**, where *volumenumber* is the volume ID obtained in the return from Step 3, and press the **Return/Enter** key.

- The volume is marked **O** (online) in the database.

11 To activate the volume, type the command **/usr/amass/bin/volstat -a <volumenumber>**, where *volumenumber* is the volume ID obtained in the return from Step 3, and press the **Return/Enter** key.

- The volume is marked **A** (active) in the database.

12 Using the report provided by STK, determine which files have had data blocks replaced with filler and delete those files from AMASS using standard UNIX commands. All such files must be recorded on a list of non-recovered files. To delete a file, type the command **rm filepathname/filename** where *filepathname/filename* is the path and name of the file, and press the **Return/Enter** key.

- The file is removed.

13 To assess dead space on the tape, type **/usr/amass/bin/vollist <volumenumber>**, where *volumenumber* is the volume ID, and press the **Return/Enter** key.

- AMASS returns information about the requested volume similar to the following example:

VOL NUM	VOL GRP	JUKE NUM	POS	VOL LABEL	FLAGS	USED (MB)	AVAIL (MB)	DEAD (%)	ERRS
18	700	1	NET	SD2102	A	35589	0	85	0

- If the amount of dead space created on the tape exceeds the allowed threshold, the files can be copied to another volume within the volume group and the tape can be reformatted (see "Recycle a Volume" in the *AMASS System Administrator's Guide*).

14 Retrieve the file location metadata to recover the ArchiveID and any checksum for each file (see Procedure 17.6.2.5.1).

- The system design incorporates calculation of a checksum when a granule is inserted into the archive. However, calculation of checksums can be time consuming, and therefore to improve system performance checksums are only calculated for a small percentage of granules on a random basis.

15 For files with a non-zero checksum returned by SDSRV (Step 14), to validate the checksum of the recovered file type **cksum <filepathname/filename>**, where *filepathname/filename* is the path and name of the recovered file.

- The system returns one line with three parameters per input file, similar to the following example:

```
cksum :Science:MOD29:2498:1.EOSHDF
1295913534 10892630 :Science:MOD29:2498:1.EOSHDF
```


The first parameter is the checksum (the second is the number of octets, and the third is the filename). If this returned checksum does not match the SDSRV-generated checksum (from Step 14), repeat Step 12 to delete the file.

- 16 For files with a checksum of zero returned by SDSRV (Step 14), it may be possible to have a Science Data Specialist use a viewing tool (e.g., EOSView) to exercise some validation on the files. It may also be possible to use information supplied by STK identifying corrupt blocks on the tape, in conjunction with the data in the **start_block_listing_volnumber** file, to determine specific files that are corrupt and recover the remaining files. The conservative approach is to assume that all zero-checksum files are corrupt and repeat Step 12 to delete them.
- 17 Add non-recovered files to the list of files to be recovered by other means (see Procedure 17.6.2.4).

Table 17.6- 6. Manual Data Recovery from Damaged Cartridge

Step	What to Do	Action to Take
1	Log in at FSMS host as amass or root	enter text; press Return/Enter
2	amassstat -c	enter text; press Return/Enter
3	/usr/amass/utlis/fileprint <filepathname/filename>	enter text; press Return/Enter
4	Remove the volume from the archive	use Procedure 17.2.4
5	Inspect the cartridge and tape for damage	observe
6	EcDsStFilesPerTapeUtility	enter text; press Return/Enter
7	Volume ID (from Step 3)	enter text; press Return/Enter
8	Send volume and file information to STK	
9	bulkinlet <volgrp> volume returned from STK	enter text; press Return/Enter
10	volloc -n <volumenumber>	enter text; press Return/Enter
11	volstat -a <volumenumber>	enter text; press Return/Enter
12	rm filepathname/filename	enter text; press Return/Enter
13	vollist <volumenumber>	enter text; press Return/Enter
14	Retrieve file location metadata	use Procedure 17.6.2.5.1
15	cksum <filepathname/filename> (files with non-zero checksum from SDSRV); delete files with non-match)	enter text; press Return/Enter
16	Other validation for files with zero checksum (e.g., EOSVIEW, recover non-corrupt files), or delete	
17	Add non-recovered files to list of files to be recovered by other means	see Procedure 17.6.2.4

17.6.2.4 Data Recovery for Known Files Not Backed Up in ECS

For any set of known files to be recovered (e.g., a list of files that were on a damaged tape and could not be recovered by Procedure 17.6.2.3, and were not available in local or offsite backups), data recovery can be attempted through procedures such as re-ingest, obtaining the data from

another DAAC that was the original source, or regeneration. Each of these potential recovery approaches is addressed in a separate procedure.

The results of file or granule recovery are slightly different depending on whether the lost files are recovered from backup, or the corresponding lost granule had to be re-archived after re-ingest or re-generation by PDPS. Files which are recovered within the Archive/STMGT procedures are re-archived under the same name, so that the affected granule(s) are restored as they were before the failure.

Where file recovery within STMGT control is not possible, granule recovery through re-ingest or re-generation results in the insertion of a new granule. This new granule has a new Universal Reference (UR) and a new 'Production Date and Time'. Particularly where granule re-generation is required, exact re-production of the original granule (data byte-for-byte) is not guaranteed.

17.6.2.4.1 Re-Ingest of Lost Data

Table 17.6-7 presents the steps required for re-ingest of lost data. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 36 Identify the source for each of the lost granules that were ingested.
- 37 If you have not already done so, retrieve the file location metadata for each file (see Procedure 17.6.2.5.1).
- 38 With reference to the applicable Interface Control Document (ICD) and using the granule metadata retrieved in Step 2, initiate the required data re-supply requests as defined in the ICD for those data suppliers able to re-supply data.
 - **Note:** Some data suppliers (e.g., Landsat-7) do not support re-supply of data.
- 39 Use appropriate ingest procedures to complete the re-ingest process.
 - Ingest procedures are addressed in Chapter 16.

Table 17.6-7. Re-Ingest of Lost Data

Step	What to Do	Action to Take
1	Identify the source for each lost granule	
2	Retrieve file location metadata	use Procedure 17.6.2.5.1
3	Initiate data re-supply with data provider	read ICD
4	Execute ingest procedure	see Chapter 16

17.6.2.4.2 Recovery of Lost Data by Reprocessing

Table 17.6-8 presents the steps required for recovery of lost data by reprocessing. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 40 If you have not already done so, retrieve the granule Production History (PH) metadata for each file for the lost granules (see Procedure 17.6.2.5.2).
- 41 Pass the output of Step 1 to the procedure "Re-Generate Granules Affected by Loss of Files from the Archive" (see Chapter 13).

Table 17.6-8. Recovery of Lost Data by Reprocessing

Step	What to Do	Action to Take
1	Retrieve granule Production History metadata	use Procedure 17.6.2.5.2
2	Execute procedure to re-generate granules affected by loss of files from the archive	see Chapter 13

17.6.2.4.2 Recovering Granules from Another Producing Site

A special case for data recovery involves granules archived at a DAAC other than the producing DAAC or site and (generally) not archived at the producing DAAC or site. Ultimately, the recovery involves a re-ingest (see Chapter 16), but because the granules are not archived at the producing DAAC or site, they must first be generated through reprocessing. Table 17.6-9 presents the steps required for recovering granules from another producing site. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 42 If you have not already done so, retrieve the file location metadata for each file for the lost granules (see Procedure 17.6.2.5.1).
- 43 Identify which of the lost granules were ingested from DAACs or other sites where the granules were produced but not archived.
 - These lost granules are known to the local SDSRV, but do not have an associated Production History (PH) granule (the PH granule is at the producing site).
- 44 Forward the granule metadata lists to the source DAAC or other site, where the metadata are used as input to the procedure "Re-Generate Granules Affected by Loss of Files from the Archive" (see Chapter 13).
- 45 Once the granule is re-generated, it may be inserted at the DAAC where it was lost, either through an order or through cross-DAAC ingest (see Chapter 16).

Table 17.6-9. Recovering Granules from Another Producing Site

Step	What to Do	Action to Take
1	Retrieve file location metadata	use Procedure 17.6.2.5.1
2	Identify lost granule(s) produced at another site	
3	Forward granule metadata to producing site	
4	Execute ingest procedure	see Chapter 16

17.6.2.5 SDSRV Procedures in Support of Data Recovery

There are two procedures described here that may be used to extract information from the SDSRV database to support the recovery of lost archive data. The first returns file metadata including file checksums for use with file recovery from backup tapes. The second generates granule metadata for use by the Planning subsystem (PLS) in re-creating granules from which files have been irrecoverably lost. The output “lists” from these procedures should be exchanged as electronic files (e.g., as email attachments) to facilitate subsequent use (e.g., to permit copying into input screens of GUIs for other procedures).

17.6.2.5.1 SDSRV Retrieval of File Location Metadata

The input to this procedure is a list of the unique file names of files in the archive affected by a tape failure (e.g., procedure 17.6.2.3 shows how the script **EcDsStFilesPerTapeUtility** may be used to generate file names for the list). The list may be called the *Affected File List* (AFL). The file names in the AFL will match the **internalFileName** column of the **DsMdFileStorage** table within the SDSRV metadata database.

The output of this procedure is a list of file metadata (archiveIDs and checksums) for each file named in the input. It may be called the *Affected File Metadata* (AFM) list. It is used to determine the backup locations, if any, of lost files and to verify the checksum of files restored through support of the tape drive vendor (StorageTek).

This procedure has the following dependencies:

- The operator is working on a machine from which SQL connections can be made to the SDSRV SQL server (e.g., e0acg11, g0acg01, l0acg02, n0acg01) and that server recognizes the Sybase account **sdsrv_role**.
- The UNIX account in use has execute permission on the required scripts, the ‘path’ shell variable set to include a directory where the command ‘isql’ is located and the SYBASE (Sybase ‘home’) environment variable set appropriately (e.g. setenv SYBASE /tools/sybOCv11.1.0).
- The operator knows the password for the SDSRV Sybase user **sdsrv_role**

Table 17.6-10 presents the steps required for SDSRV retrieval of file location metadata. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 46** Log in at the SDSRV host (e.g., e0acs05, g0acs03, l0acs03, n0acs04).
- 47** Receive the Affected File List (AFL) (e.g., file list from output of script **EcDsStFilesPerTapeUtility**) as an electronic file; save a local copy of the file with an identifiable name (e.g., aflfile.txt).
- 48** To change directory to the location of the database scripts, type **cd /usr/ecs/<MODE>/CUSTOM/dbms/DSS** and press the **Return/Enter** key.
 - The **<MODE>** will most likely be one of the following operating modes:
 - OPS (for normal operation).

– TS1 or TS2 (for testing).

- The working directory is changed to `/usr/ecs/<MODE>/CUSTOM/dbms/DSS`.

4 9 To execute the script for retrieval of file location metadata, type **DsDbSrFileLocMetadata** *aflfile.txt aflmetadata.txt*, where *aflfile.txt* is the name of the input file with the list of affected files and *aflmetadata.txt* is the desired name of the output file, and then press the **Return/Enter** key.

- The script requires that certain environmental variables be set prior to execution. If you have not set them, the script returns an error message listing the variables that must be set and giving examples. To make the script execute properly, you may need to set environmental variables using the following commands and appropriate variable entries:

```
setenv DSQUERY x0acg0n_srvr (e.g., e0acg11, g0acg01, l0acg02, n0acg01)
```

```
setenv DBNAME EcDsScienceDataServer1_<MODE>
```

```
setenv DBUSERNAME sdsrv_role
```

```
setenv DBPASSWD <password>
```

```
setenv SYBASE /tools/sybOCv11.1.1
```

- If the variables are set appropriately, the script uses the data in the input file to generate the named output file; during the execution it provides feedback similar to the following:

```
Using Login   : sdsrv_role
Using Server  : tlacg04_srvr
using Database: EcDsScienceDataServer1_TS2
```

```
Recovering the Effected Lost Files....
```

```
**** No errors found in DBoutfile_FileLocMetadata ***
```

50 To check that the output file is not empty (i.e., of zero length), type the command **ls -l aflmetadata.txt**, where *aflmetadata.txt* is the name of the output file, and press the **Return/Enter** key.

- UNIX displays the file information in the following form:

```
-rw-r--r--  1 cmshared      293 Sep 25 15:25 aflmetadata.txt
```

- If the file is of zero length, either the input file was of zero length or an unexplained error occurred. Check the input file.

51 To visually inspect the file to verify the success of the command, type the command **view aflmetadata.txt**, where *aflmetadata.txt* is the name of the output file, and press the **Return/Enter** key.

- The contents of the output file are in two sections: 1) the affected file metadata found within SDSRV inventory database; and 2) files not found within SDSRV database (this section is usually empty). UNIX displays the contents of the file in the following form (for this sample data, the input file contained only one filename):

```
InternalFileName_found_in_SDSRV_Inventory_Metadata_Database
Tue Sep 25 15:25:49 EDT 2001
```

```
:SC:MOD000.001:19862:1.CCSDS
"Aug 21 2001 3:34:17:000PM" 0 1670000364 SCMOD000.00119862
```

```
InternalFileName_not_found_in_SDSRV_Inventory_Metadata_Database
Tue Sep 25 15:25:50 EDT 2001
```

- To exit the view process, type **:q!** and press the **Return/Enter** key. (*Note:* This step specifies use of the **view** command to view the file, but the content can be viewed using other commands as well [e.g., **vi**, **pg**, **more**]).

52 It is advisable, especially if there are large numbers of affected files, to check for errors in the output of the script, searching for occurrences of the strings 'msg' and 'error.' To execute a check for 'msg,' type the command **grep -i msg aflmetadata.txt | wc -l**, where *aflmetadata.txt* is the name of the output file, and press the **Return/Enter** key. To execute a check for 'error,' type the command **grep -i error aflmetadata.txt | wc -l** and press the **Return/Enter** key.

- If no errors occurred, UNIX returns an output of '0' (zero).
- Any other output means that there were errors in the process. If errors are found, they must be diagnosed based on the error message(s) and the procedure must be repeated after correction of the input file.

53 When the output file passes the tests of Step 7, it can be passed to the calling procedure.

Table 17.6-10. SDSRV Retrieval of File Location Metadata

Step	What to Do	Action to Take
1	Log in at SDSRV host	enter text; press Return/Enter
2	Save a local copy of the Affected File List	name file, e.g., aflfile.txt
3	cd /usr/ecs/<MODE>/CUSTOM/dbms/DSS	enter text; press Return/Enter
4	DsDbSrFileLocMetadata aflfile.txt aflmetadata.txt	enter text; press Return/Enter
5	ls -l aflmetadata.txt	enter text; press Return/Enter ; read text
6	view aflmetadata.txt	enter text; press Return/Enter ; read text
7	grep -i msg aflmetadata.txt wc -l grep -i error aflmetadata.txt wc -l	enter text; press Return/Enter ; read text
8	Pass output file to calling procedure	

17.6.2.5.2 SDSRV Retrieval of Granule Production History Metadata

The input to this procedure is a list of files remaining to be recovered after as many files as possible are recovered from backup (see Procedures 17.6.2.1 and 17.6.2.2) and/or from a damaged cartridge (see Procedure 17.6.2.3). The list may be referred to as the *Affected File List* (AFL). The output of this procedure serves as input to the PDPS/PLS procedure for granule regeneration, “Re-Generate Granules Affected by Loss of Files from the Archive,” in Chapter 13.

The goal of this procedure is to list PLS-required granule metadata for use by the local PDPS in re-generating lost granules. The procedure extracts valid Production History URs for the ‘lost’ granules from the local SDSRV database. It assumes that the Delivered Algorithm Package information has been inserted into the SDSRV for all associate datatypes in the process of Science Software Integration and Test (SSI&T), which is addressed in Chapter 26. There may be a period in the lifetime of a granule when this information is not populated, during which granule attributes PGENAME and PGEVersion are not available to this procedure.

Output from this procedure is a file containing:

- “Granules for PDPS Re-generation” -- those found within SDSRV. These are passed to the PLS operators for re-generation using the procedure “Re-Generate Granules Affected by Loss of Files from the Archive” (see Chapter 13).
- “Residual Granules to Recover” -- those not found within SDSRV inventory. This list may include granules that have been removed by the process for physical deletion from the archive and SDSRV inventory (see Procedure 17.4.3).

Table 17.6-11 presents the steps required for SDSRV retrieval of granule production history metadata. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 54** Log in at the SDSRV host (e.g., e0acs05, g0acs03, l0acs03, n0acs04).
- 5 5** Receive the Affected File List (AFL) (e.g., file list from output of script **EcDsStFilesPerTapeUtility**) as an electronic file; save a local copy of the file with an identifiable name (e.g., aflfile.txt).
- 56** To change directory to the location of the database scripts, type **cd /usr/ecs/<MODE>/CUSTOM/dbms/DSS** and press the **Return/Enter** key.
- The **<MODE>** will most likely be one of the following operating modes:
 - OPS (for normal operation).
 - TS1 or TS2 (for testing).
 - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/dbms/DSS**.
- 5 7** To execute the script for retrieval of file location metadata, type **DsDbSrGranPHMetadata aflfile.txt agrmetadata.txt**, where **aflfile.txt** is the name of the input file with the list of affected files and **agrmetadata.txt** is the desired name of the output file, and then press the **Return/Enter** key.

- The script requires that certain environmental variables be set prior to execution. If you have not set them, the script returns an error message listing the variables that must be set and giving examples. To make the script execute properly, you may need to set environmental variables using the following commands and appropriate variable entries:

```
setenv DSQUERY x0acgnn_srvr (e.g., e0acg11, g0acg01, l0acg02, n0acg01)
```

```
setenv DBNAME EcDsScienceDataServer1_<MODE>
```

```
setenv DBUSERNAME sdsrv_role
```

```
setenv DBPASSWORD <password>
```

```
setenv SYBASE /tools/sybOCv11.1.1
```

- If the variables are set appropriately, the script uses the data in the input file to generate the named output file; during the execution it provides feedback similar to the following:

```
Using Login   : sdsrv_role
Using Server  : t1acg04_srvr
using Database: EcDsScienceDataServer1_TS2
```

```
Recovering the Lost Files....
```

```
**** No errors found in DBoutfile_GranPHMetadata ***
```


58 To check that the output file is not empty (i.e., of zero length), type the command **ls -l agrmetadata.txt**, where *agrmetadata.txt* is the name of the output file, and press the **Return/Enter** key.

- UNIX displays the file information in the following form:

```
-rw-r--r--  1 cmshared      293 Sep 25 15:25 agrmetadata.txt
```

- If the file is of zero length, either the input file was of zero length or an unexplained error occurred. Check the input file.

59 To visually inspect the file to verify the success of the command, type the command **view agrmetadata.txt**, where *agrmetadata.txt* is the name of the output file, and press the **Return/Enter** key.

- The contents of the output file are in two sections: 1) the granule metadata found within SDSRV inventory database; and 2) granule metadata not found within SDSRV database (residual granules to recover). For each of the files listed in the input file for which related granule metadata are found in the SDSRV, the script output should include the GeoID (partial UR), the UR of any available associated Production History granule, the ESDT shortname and versionID, the granule beginning date and time and ending date and time. UNIX displays the contents of the file in the following form:

```
Granule_metadata_found_within_SDSRV_Inventory_database Wed Sep 26
11:11:51 EDT 2001
```

```
:BR:Browse.001:1170:1.BINARY
```

```
0 1000 BRBrowse.0011170 PGName 1          "None"  "None"  "NONE"
"NORMAL"
```

```
PH_Does_Not_Apply
```

```
:SC:MOD00:65001:1.CCSDS
```

```
0 1000 SCAST_04.00120001 PGName 1
"Jan  1 1997 12:00:00:000AM" "Jan  1 1997 12:00:00:000AM" "Oct 10
1996 12:02:00:000AM" "NORMAL" 2
```

```
NO_PH
```

```
:PH:PH.001:2000000076:1.BINARY
```

```
0 65536 PHPH.0012000000076 PGName 1          "None"  "None"
"None"  "NORMAL"
```

```
PH_Does_Not_Apply
```

```
:QA:QA.001:1003:1.ASCII
```

```
0 0 QAQA.0011003 PGName 1          "None"  "None"  "None"
"NORMAL"
```

```
PH_Does_Not_Apply
```

```
:OR:OR.001:2102:1.ASCII
```

```
0 0 OROR.0012102 PGName 1 "None" "None" "None"  
"NORMAL"
```

```
PH_Does_Not_Apply
```

```
:AN:AN.001:3100:1.ASCII
```

```
0 0 ANAN.0013100 PGName 1 "None" "None" "None"  
"NORMAL"
```

```
PH_Does_Not_Apply
```

```
Granule_metadata_not_found_within_SDSRV_Inventory_database Wed Sep  
26 11:11:51 EDT 2001
```

- To exit the view process, type **:q!** and press the **Return/Enter** key. (*Note:* This step specifies use of the **view** command to view the file, but the content can be viewed using other commands as well [e.g., **vi**, **pg**, **more**]).

60 It is advisable, especially if there are large numbers of affected files, to check for errors in the output of the script, searching for occurrences of the strings 'msg' and 'error.' To execute a check for 'msg,' type the command **grep -i msg agrmetadata.txt | wc -l**, where *agrmetadata.txt* is the name of the output file, and press the **Return/Enter** key. To execute a check for 'error,' type the command **grep -i error aflmetadata.txt | wc -l** and press the **Return/Enter** key.

- If no errors occurred, UNIX returns an output of '0' (zero).
- Any other output means that there were errors in the process. If errors are found, they must be diagnosed based on the error message(s) and the procedure must be repeated after correction of the input file.

61 When the output file passes the tests of Step 7, it can be passed to the calling procedure (e.g., Chapter 13 Procedure to Re-Generate Granules Affected by Loss of Files from the Archive).

- *Note:* Granules for recovered files will, by definition, have a different granuleURs (dbIDs) than the files that were lost.

Table 17.6-11. SDSRV Retrieval of Granule Production History Metadata

Step	What to Do	Action to Take
1	Log in at SDSRV host	enter text; press Return/Enter
2	Save a local copy of the Affected File List	name file, e.g., <i>afffile.txt</i>
3	cd /usr/ecs/<MODE>/CUSTOM/dbms/DSS	enter text; press Return/Enter
4	DsDbSrGranPHMetadata <i>afffile.txt agrmetadata.txt</i>	enter text; press Return/Enter
5	ls -l <i>agrmetadata.txt</i>	enter text; press Return/Enter ; read text
6	view <i>agrmetadata.txt</i>	enter text; press Return/Enter ; read text
7	grep -i msg <i>agrmetadata.txt</i> wc -l grep -i error <i>agrmetadata.txt</i> wc -l	enter text; press Return/Enter ; read text
8	Pass output file to calling procedure	

17.7 Archive Troubleshooting

There are several troubleshooting tools provided with AMASS that can assist you in monitoring archive activity and in responding to fault notifications. The *AMASS System Administrator's Guide* (available electronically on **drg** servers [e.g., g0drg01, e0drg11, l0drg01, n0drg01] in directory **/usr/amass/books**) includes instructions on using these tools. Some of the most useful ones are addressed in this section. Table 17.7-1 provides an Activity Checklist for archive troubleshooting.

Table 17.7-1. Activity Checklist for Archive Troubleshooting

Order	Role	Task	Section	Complete?
1	Archive Manager	Checking daemons and using <i>healthcheck</i>	(P) 17.7.1.1	
2	Archive Manager	Using <i>sysperf</i> to Display the Status of AMASS I/O Activity	(P) 17.7.1.2	
3	Archive Manager	Using <i>vollist</i> to Display Volume Data	(P) 17.7.1.3	
4	Archive Manager	Using the <i>amass_log</i> Script to Display AMASS Errors	(P) 17.7.1.4	
5	Archive Manager	Using <i>quedisplay</i> to View What is in the AMASS Queue	(P) 17.7.1.5	
6	Archive Manager	Using <i>mediamove</i> to Establish Synchrony Between <i>quedisplay</i> and <i>medialist</i>	(P) 17.7.1.6	
7	Archive Manager	Checking Log Files	(P) 17.7.2.1	
8	Archive Manager	A Special Case: Checking the Request Manager Server Debug Log	(P) 17.7.2.2	
9	Archive Manager	Checking the el_ETAC Log File	(P) 17.7.2.3	
10	Archive Manager/ Database Administrator	Handling a Data Insertion Failure	(P) 17.7.2.4	
11	Archive Manager	Handling a Data Acquire Failure	(P) 17.7.2.5	
12	Archive Manager	Diagnosing/Investigating Write Errors	(P) 17.7.3	
13	Archive Manager	Diagnosing/Investigating Read Errors	(P) 17.7.4	

17.7.1 Using AMASS Commands, Utilities, and Scripts for Monitoring and Fault Response

The AMASS file system needs to have the following daemons running at all times:

- **amassmain.**
- **daemons/lm_ip -a fslock.**
- **qset.**
- **klogd.**
- **amass_iocomp.**
- **libsched** (one instance for each virtual library).
- **libio_tape** (at least one instance for each drive in each jukebox).

The UNIX process search provides an easy check for these daemons. If they are up, the AMASS **healthcheck** command provides a useful check on the health of AMASS while it is running.

A command provided to display the status of the AMASS I/O activity is **sysperf**. This command returns several items:

- the number of reads and writes that are outstanding.
- the number of volumes (for reads) or volume groups (for writes) that are going to be used by those reads and writes.
- the current volumes in the drives.
- the I/O rate in Kb per second since the last update. This value first appears as a zero. Then AMASS continues to update the information at intervals based on a value for *updateinterval* entered by the operator.

Sysperf can often show the first sign of trouble. For example, if there are reads and writes in process but throughput is always 0, a problem is indicated. The most common problems are volumes and drives that go off line and/or inactive.

Volumes are monitored using the **vollist** command. It returns information on the status of a specified volume or list of volumes in the archive. If the output of **vollist** indicates that the volume is inactivated (i.e., **I** appears in the **FLAGS** column), the **amass_log** script can help to determine the nature of the problem. The **amass_log** script displays AMASS messages from the system log file. This script can provide helpful information under several circumstances, such as when a command gives unexpected results or when AMASS appears not to be functioning properly in other ways.

Unless use of the **amass_log** script shows that there are many errors on a volume that has been inactivated, you can reactivate the volume using the command:

```
/usr/amass/bin/volstat -a <volumenumber>
```

where **<volumenumber>** is the volume ID for the volume to be activated.

Just as **vollist** provides information on the status of volumes, the command **drivelist** displays the status of drives available to AMASS. Active drives are noted by an **A**, and inactive drives are noted by an **I**. The command is **/usr/amass/bin/drivelist**. If AMASS inactivates a drive, use the **amass_log** script as described previously. Unless there is a hardware problem and several attempts have been made to ready the drive, it is usually appropriate to reactivate the drive using the **drivestat** command. For example, to reactivate drive 1 in jukebox 1, type the command **/usr/amass/bin/drivestat -a 1 1**.

A useful library utility included with AMASS is **quedisplay**. This utility permits the operator to see what is in the queue, and to diagnose problems such as the following:

- During an attempt to write to a file, the drive light does not illuminate.
- The system is slowing down.
- An AMASS command does not complete.

The output of the **quedisplay** utility shows the queue, which consists of read and write requests, AMASS administration commands, and a list of libraries, drives, and what volumes they manage. An example of output from this utility might take the following form:

```
READQ rid=52696, fptr=0xf0227c5c, vol=3, fnode_flags=0x110
WRITEQ rid=79, fptr=0xc00eff54, vol=5, fnode_flags=0x8048844
ADMINQ:cmd=1, flags=0x6,vol=32, juke=1, pid=1047, ftype=0, err=0
JUKEBOX 1 DRIVE 1, vid=32, vflag=0x100, status=0
JUKEBOX 1 DRIVE 2, no volume in drive
```

If there are **READQ** or **WRITEQ** entries, the name(s) of the file(s) being processed can be determined by using the **filepath** command and the first number in the entry. For example, enter **/usr/amass/utls/filepath 52696** for the first file number in the sample output.

Occasionally, a robot may lose synchrony with AMASS concerning the location of media. The best way to verify this is to compare **quedisplay** and **medialist**. The **medialist** utility is a standalone program that communicates with the robot controller in the Library Storage Module to determine the robot's view of media and their slot locations. If the two programs disagree, you can bring the two programs into synchrony using **mediamove**. The following paragraphs provide step-by-step procedures for use of some of these commands and utilities.

17.7.1.1 Checking Daemons and Using *healthcheck*

If there is an indication or question of a potential problem with AMASS, an appropriate initial step is to check the status of the required daemons. If the check indicates that the daemons are up, then it is a reasonable next step to run the *healthcheck* command. Table 17.7-2 presents the steps required for checking daemons and using *healthcheck* to verify the status of AMASS. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in as **amass** at the FSMS host.
- 2 Type **ps -ef | grep amass** and press the **Return/Enter** key.
 - UNIX returns running AMASS processes in a format similar to the following:

```

amass 7214464 7208385 0 Sep 19 ? 0:00 libio_tape 2 1
amass 7208385 1 0 Sep 19 ? 10:33 /usr/amass/daemons/amassmain 0
amass 7214747 7208385 0 Sep 19 ? 0:10 amass_iocomp
amass 7282853 7215637 0 Sep 20 ? 1:47 libio_tape 1 1
amass 7282868 7215637 0 Sep 20 ? 0:00 libio_tape 1 1
amass 6949087 7215637 0 Sep 20 ? 1:47 libio_tape 1 1
amass 7214915 7208385 0 Sep 19 ? 0:00 klogd
amass 7214972 7208385 0 Sep 19 ? 50:54 libio_tape 1 2
amass 5539722 7217884 0 Sep 20 ? 0:23 libio_tape 1 3
amass 7301726 7215964 0 Sep 20 ? 1:05 libio_tape 3 1
amass 7215313 1 0 Sep 19 ? 9:34 /usr/amass/daemons/lm_ip -a
fslock1 -u 128 -f 256 -q 128
amass 7357656 7216363 0 Sep 20 ? 0:00 libio_tape 3 3
amass 7215637 7208385 0 Sep 19 ? 84:10 libio_tape 1 1
amass 7215638 7208385 0 Sep 19 ? 2:43 libsched 3
amass 7277545 7214972 0 Sep 20 ? 0:41 libio_tape 1 2
amass 7215870 7208385 0 Sep 19 ? 2:52 libsched 1
amass 7215964 7208385 0 Sep 19 ? 109:25 libio_tape 3 1
amass 7216363 7208385 0 Sep 19 ? 84:16 libio_tape 3 3
amass 6950984 7217884 0 Sep 20 ? 0:23 libio_tape 1 3
amass 8175053 7212410 0 Sep 26 ? 0:00 libio_tape 1 4
amass 7340525 7217134 0 Sep 20 ? 1:19 libio_tape 3 2
amass 7278745 7217884 0 Sep 20 ? 0:23 libio_tape 1 3
amass 7216941 7208385 0 Sep 19 ? 0:32 qset
amass 7340710 7217134 0 Sep 20 ? 1:19 libio_tape 3 2
amass 7217134 7208385 0 Sep 19 ? 138:26 libio_tape 3 2
amass 7359550 7216363 0 Sep 20 ? 0:52 libio_tape 3 3
amass 7217388 7208385 0 Sep 19 ? 0:00 libio_tape 2 2
amass 7285477 7215637 0 Sep 20 ? 1:47 libio_tape 1 1
amass 7285537 7215637 0 Sep 20 ? 1:47 libio_tape 1 1
amass 7217884 7208385 0 Sep 19 ? 17:37 libio_tape 1 3
amass 7279821 7214972 0 Sep 20 ? 0:41 libio_tape 1 2
amass 6878049 7208385 0 Sep 19 ? 2:36 libsched 2
amass 7279907 7214972 0 Sep 20 ? 0:41 libio_tape 1 2
amass 7335573 7217134 0 Sep 20 ? 1:19 libio_tape 3 2
.
.
.

```

- If the running processes do not include **amassmain**, **daemons/lm_ip -a fslock**, **qset**, **klogd**, **amass_iocomp**, **libsched**, and **libio_tape**, it may be necessary to restart AMASS (refer to Procedure 17.1.3 **Rebooting AMASS**).

- 3 To check the AMASS database integrity, check the availability of write resources FNODEs and cache blocks, and to verify cache partitions, type **/usr/amass/bin/healthcheck -viwc** and press the **Return/Enter** key.

- AMASS returns information on its health in format similar to the following:

```
--- STARTING DATABASE INTEGRITY CHECK ---

    -api has been opened properly and AMASS is running.
    -verifying pathnames.
    -got locks on database
    -unlocking database tables and exiting

--- CHECK COMPLETED!! ---

--- CHECKING AVAILABILITY OF WRITE RESOURCES FNODEs AND CACHE
BLOCKS ---

    -api has been opened properly and AMASS is running.
    -Initializing the passed arguments.
    -Returning the passed arguments.
    -Restoring signals.
    -exiting.

--- CHECK COMPLETED!! ---

--- RUNNING CACHE TEST ---

    -api has been opened properly and AMASS is running.
    -Validating the raw cache.
    -Restoring signals.
    -exiting.

--- TEST COMPLETED!! ---
```

- If an error message is returned, it may be necessary to restart AMASS (refer to Procedure 17.1.3 **Rebooting AMASS**).

- 4 To check library components, type `/usr/amass/bin/healthcheck -vl 1 0 volumenumber` and press the **Return/Enter** key.
- The argument **-l** (lower-case l) specifies the library components check, and requires specification of a jukebox (**1** in this case), a drive number (entering **0** as in this case checks all active drives), and a volume number (*volumenumber* is the volume ID of an available volume in the specified jukebox; it may be helpful to use the **vollist** command [refer to Procedure 17.7.1.3 **Using vollist to Display Volume Data**] to identify a suitable volume, such as a volume in the Space Pool, to use for this test).
 - AMASS returns information on the health of library components in the following format:

```

--- CHECKING LIBRARY COMPONENTS ---

-api has been opened properly and AMASS is running.
-mapping shared memory.
-verifying the juke number.
-validating volume number.
-validating drive number and checking for active drive/s/.
-saving the volume's status before inactivating it.
-proceeding with physical test.
-restoring signals and exiting.

--- CHECK COMPLETED!! ---

```
 - If an error message is returned, it may be necessary to restart AMASS (refer to Procedure 17.1.3 **Rebooting AMASS**) and/or to check for possible hardware problems with drives or other components.

Table 17.7-2. Checking Daemons and Using healthcheck

Step	What to Do	Action to Take
1	Log in at FSMS host	enter text; press Return/Enter
2	ps -ef grep amass	enter text; press Return/Enter
3	/usr/amass/bin/healthcheck -viwc	enter text; press Return/Enter
4	/usr/amass/bin/healthcheck -vl 1 0 volumenumber	enter text; press Return/Enter

17.7.1.2 Using *sysperf* to Display the Status of AMASS I/O Activity

Table 17.7-3 presents the steps required for using *sysperf* to display the status of AMASS I/O activity. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01).
- 2 Type **/usr/amass/bin/sysperf -k 5** and press the **Return/Enter** key.
 - The screen updates every 5 seconds and display information on the amass kernel (**-k**) in a form similar to the following example (**Note**: A different number of seconds may be entered to specify a different refresh rate.):


```

SYSTEM STATISTICS - Thu Sep 27 08:17:33
UPDATE INTERVAL   - 10 SEC
AVERAGE THROUGHPUT - 0 KBYTES/SEC

READ REQUESTS      # OF VOLUMES
0                  0

WRITE REQUESTS     # OF VOL GROUPS
0                  0

CACHE BLOCKS      2957 Total    2957 Free      0 Dirty
FNODES            800 Total    796 Free      4 Used

JUKE  DRIVE  VOLFLAGS  VOLUME  VOLGRP  KBYTES/SEC
          
```
- 3 To break out of the command, use **ctrl-c** (while holding down the **Control Key**, press **c**).
 - The screen stops updating and displays a UNIX prompt.

Table 17.7-3. Using sysperf to Display the Status of AMASS I/O Activity

Step	What to Do	Action to Take
1	Log in at FSMS host	enter text; press Return/Enter
2	sysperf -k 5	enter text; press Return/Enter
3	ctrl-c	hold Control key and press c

17.7.1.3 Using *vollist* to Display Volume Data

Table 17.7-4 presents the steps required for using *vollist* to display volume data. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 5 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01).
- 6 Type **/usr/amass/bin/vollist *volumenumber*** (where *volumenumber* is the ID for one of the volumes) and press the **Return/Enter** key.
 - AMASS displays volume data in the following form:

```

VOL  VOL  JUKE POS  VOL      FLAGS  USED  AVAIL  DEAD  ERRS
NUM  GRP  NUM          LABEL      (MB)  (MB)  (%)
100  500  3    NET  SD0060  O      99213  3167   0     0

```

- **Note:** In this example for volume 100, the **O** in the **FLAGS** column indicates that the volume is offline. Other often-used flags are: **A** for Active, **I** for Inactive, **R** for Read-only, **U** for Unformatted.
 - If *volumenumber* is omitted from the command, AMASS displays volume data for all volumes.
 - If the argument **-g** is used with the command and a volume group identifier is specified (i.e., **vollist -g volumegroupnumber**), AMASS displays volume data for each volume in the specified volume group.
- 7 To put an offline volume back on line, type **/usr/amass/bin/volloc -n volumenumber** and press the **Return/Enter** key.
- The specified volume is put online, and in output from execution of an appropriate **vollist** command, AMASS displays **A** in the **FLAGS** column.

Table 17.7-4. Using vollist to Display Volume Data

Step	What to Do	Action to Take
1	Log in at FSMS host	enter text; press Return/Enter
2	vollist [-g] [volumenumber]	enter text; press Return/Enter
3	volloc -n volumenumber (to put offline volume online)	enter text; press Return/Enter

17.7.1.4 Using the *amass_log* Script to Display AMASS Errors

Table 17.7-5 presents the steps required for using the *amass_log* script to display AMASS errors. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 8 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01).
- 9 To change to the AMASS tools directory, type **cd /usr/amass/tools**, and then press the **Return/Enter** key.
 - The working directory is changed to **/usr/amass/tools**.
- 10 Type **./amass_log logfilepath**, where *logfilepath* is the full pathname of the system log file to scan for AMASS messages, and then press the **Return/Enter** key.
 - On a Sun, the *logfilepath* is likely to be **/var/adm/messages**; on an SGI, the *logfilepath* is likely to be **/var/adm/SYSLOG**. Any AMASS error messages in the scanned log file are displayed.
- 11 Perform the action recommended for the error message in the log.
 - The *AMASS System Administrator's Guide* (available electronically on **drg** servers [e.g., g0drg01, e0drg11, l0drg01, n0drg01] in directory **/usr/amass/books**) provides detailed information concerning error messages. An error message informs of critical problems that prevent AMASS from functioning. An error message is usually followed by a correction message, which provides instructions for correcting the situation. Sometimes, there is a previous warning message that may provide an accompanying correction message. Other messages may be identified by number only; the *System Administrator's Guide* provides a reference list, with accompanying corrective actions.

Table 17.7-5. Using the amass_log Script to Display AMASS Errors

Step	What to Do	Action to Take
1	Log in at FSMS host	enter text; press Return/Enter
2	cd /usr/amass/tools	enter text; press Return/Enter
3	./amass_log logfilepath	enter text; press Return/Enter
4	Perform recommended action (see <i>System Administrator's Guide</i>)	read text

17.7.1.5 Using *quedisplay* to View What is in the AMASS Queue

Table 17.7-6 presents the steps required for using *quedisplay* to view what is in the AMASS queue. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 12 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01).
- 13 To change to the utilities directory, type **cd /usr/amass/utills**, and then press the **Return/Enter** key.

- The working directory is changed to **/usr/amass/utills**.

- 14 Type **quedisplay**, and then press the **Return/Enter** key.

- The AMASS queue is displayed in the following form:

```

READQ rid=52696, fptr=0xf0227c5c, vol=3, fnode_flags=0x110
WRITEQ rid=79, fptr=0xc00eff54, vol=5, fnode_flags=0x8048844
ADMINQ:cmd=1, flags=0x6,vol=32, juke=1, pid=1047, ftype=0, err=0
JUKEBOX 1 DRIVE 1, vid=32, vflag=0x100, status=0
JUKEBOX 1 DRIVE 2, no volume in drive

```

Note: In the output, "rid" = Record ID, "pid" = Process ID

Table 17.7-6. Using *quedisplay* to View What is in the AMASS Queue

Step	What to Do	Action to Take
1	Log in at FSMS host	enter text; press Return/Enter
2	cd /usr/amass/utills	enter text; press Return/Enter
3	quedisplay	enter text; press Return/Enter

17.7.1.6 Using *mediamove* to Establish Synchrony Between *quedisplay* and *medialist*

Table 17.7-7 presents the steps required for using *mediamove* to establish synchrony between *quedisplay* and *medialist*. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 15 Log in as **amass** at the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01).

- 16 Type **/usr/amass/utills/quedisplay** and then press the **Return/Enter** key.

- AMASS displays the following information (for example of incorrect status).

```

. . .
JUKEBOX 1 DRIVE 1, no volume in drive
JUKEBOX 1 DRIVE 2, vid=50, vflags=0x4, status=0

```

- 17 Type **/usr/amass/utills/medialist** and then press the **Return/Enter** key.

- AMASS displays the following information (for example of actual status).

```

. . .
SLOT VSD0098 FULL
DRIVE 1 FULL FROM VSD0096
DRIVE 2 FULL FROM VSD0097

```

- Note that the **medialist** result shows that drive 1 actually is occupied, although **quedisplay** registers that drive 1 is empty.

18 Type **/usr/amass/utlils/mediamove 1 VSD0096 1** and then press the **Return/Enter** key.

- AMASS moves the volume from the *source* (drive **1** in this example) to the *destination* (slot **VSD0096** in this example) in the specified *jukeboxnumber* (jukebox **1** in this example), thereby bringing the actual status of drive 1 (as known by *medialist*) to the status reflected by *quedisplay*.

**Table 17.7-7. Using mediamove
to Establish Synchrony Between quedisplay and medialist**

Step	What to Do	Action to Take
1	Log in as amass at FSMS host	enter text; press Return/Enter
2	quedisplay	enter text; press Return/Enter
3	medialist	enter text; press Return/Enter
4	mediamove source destination [jukeboxnumber]	enter text; press Return/Enter

17.7.2 Recovering from Failure to Store or Retrieve Data

Successful data storage and retrieval functions are the heart of ECS. Successful ingest of data or processing of data to produce new science data granules require that Storage Management (STMGT) is inserting the product into the archive and that Science Data Server (SDSRV) is inserting the associated metadata into the inventory. Staging disks and cache managers for the Archive server and the FTP server are also involved in this process. To check the functioning of these elements, it is necessary that the ESDTs for the data to be inserted are installed and available, and that subscriptions have been registered.

Troubleshooting failures to store or retrieve data (as well as other failures) often requires review of server or application log files. This section contains a general procedure for reviewing log files to check for proper start-up and communications. It also has a procedure for a special case of reviewing log files for the Storage Management Request Manager server, and a procedure for reviewing the current el_ETAC log file of interactions between AMASS and ACSLS. Separate procedures then address recovery from failure to insert (store) data and recovery from failure to acquire (retrieve) data.

17.7.2.1 Checking Server Log Files

Table 17.7-8 presents the general steps required for checking server log files. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1** Log in to the host for the server and log(s) to be examined.
- 2** Type **cd /usr/ecs/<MODE>/CUSTOM/logs** and then press the **Return/Enter** key.
 - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/logs**.

- 3 To view a server log, type **pg filename** and then press the **Return/Enter** key.
 - *filename* refers to the log file to be reviewed (e.g., **EcDsScienceDataServer.ALOG**, **EcDsScienceDataServerDebug.log**).
 - The first page of the log file is displayed; additional sequential pages can be displayed by pressing the **Return/Enter** key at the **:** prompt.
 - Although this procedure has been written for the **pg** command, any UNIX editor or visualizing command (e.g., **vi**, **view**, **more**, **tail**) can be used to review the log file.
 - Typically, the **<server>Debug.log** captures more detailed information than the **<server>.ALOG**. However, for some servers (e.g., **SDSRV**), there may be significant detail in the **<server>.ALOG**. It is also important to note that the **DebugLevel** parameter setting in the **Configuration Registry** determines the level of detail captured in the **<server>Debug.log** (0 is off, a setting of **1** captures status and errors, a setting of **2** captures major events, and a setting of **3** is a full trace recording of all activity). If the **DebugLevel** has been set to one of the lower levels during operations, the System Administrator may set it to **3** during troubleshooting. Similarly, the **AppLogLevel** parameter setting determines the level of detail captured in the **<server>.ALOG** (0 provides a full trace recording of all events, 1 provides messages related to all major events, 2 yields just records of errors, and 3 turns recording off). (Note: As of Release 6A, there are new debug levels available for some logs; Storage Management (STMGT) offers "enhanced" debugging based on bitmasks. Level 7 provides a four-bit level for detailed database debugging. Level 15 provides an eight-bit level that repeatedly dumps the in-memory request queue in the STMGT Request Manager.
- 4 Review the log file(s) to determine if there are any indications of connection problems or errors at start up.
 - The log file for the called server may contain an error message indicating a problem at start-up. The debug log should indicate a typical start sequence, including (sample log entries in the following material were taken from a debug log showing start-up for **EcDsStFtpServer**):
 - Get parameters from registry (log entries similar to the following).

```
DSS EcDsStFtpServer Server Debug log on f2acg01 starting at Mon
Jun  4 07:57:45 EDT 2001
EcAgInstanceID Sequence Number is 3870
Setting up environment variables needed for DCE:
RPC_UNSUPPORTED_NETIFS = ""
/usr/ecs/DEV07/CUSTOM/bin/DSS/EcDsStFtpServer ConfigFile
/usr/ecs/DEV07/CUSTOM/cfg/EcDsStFtpServer.CFG ecs_mode DEV07
StartTemperature cold
Started process EcDsStFtpServer in mode DEV07 with PID 2709893
EcRgRegistry_1_0::ctor this = 0x104eef38
EcRgRegistry_1_0::ctor this = 0x104eef88
FoIpPtToPtPortalImp::Send sent 20/20
```

```

FoIpPtToPtPortalImp::Send sent 219/219
FoIpPtToPtPortalImp::Receive got 20
FoIpPtToPtPortalImp::Receive got 1024
FoIpPtToPtPortalImp::Receive got 246
***** After Retrieving of RGY: Name = EcDsStFtpServerNONE
ProgramID = 4645102
ApplicationID = 4600000
Release = B
DeltaTime = 0
Site = RBD
SubSysName = DSS
MajorVersion = 1
MinorVersion = 0
DebugLevel = 3
AppLogLevel = 0
AppLogSize = 3000000
DBServer = f2acg01_srvr
DBLoginName = EcDsStFtpServer
DBName = stmgtdbl

```

- **Load resource catalogs (log entries indicate the loading, or that the loading did not complete, similar to the following).**

```

06/04/01 07:57:47: Thread ID : 65536 : loading resource catalog
file from
/usr/ecs/DEV07/CUSTOM/data/DSS/ResourceCatalogs/DsMdResource.dat.
rcat
06/04/01 07:57:48: Thread ID : 65536 : loading resource catalog
file from
/usr/ecs/DEV07/CUSTOM/data/DSS/ResourceCatalogs/EcDsSdHr.dat.rcat
06/04/01 07:57:48: Thread ID : 65536 : loading resource catalog
file from
/usr/ecs/DEV07/CUSTOM/data/DSS/ResourceCatalogs/DsSrResource.dat.
rcat
06/04/01 07:57:48: Thread ID : 65536 : loading resource catalog
file from
/usr/ecs/DEV07/CUSTOM/data/DSS/ResourceCatalogs/DsGlResource.dat.
rcat
06/04/01 07:57:48: Thread ID : 65536 : loading resource catalog
file from
/usr/ecs/DEV07/CUSTOM/data/DSS/ResourceCatalogs/DsShResource.dat.
rcat
06/04/01 07:57:48: Thread ID : 65536 : loading resource catalog
file from
/usr/ecs/DEV07/CUSTOM/data/DSS/ResourceCatalogs/EcDsSdHc.dat.rcat

```

- **Pre-cache errors associated with database connectivity (log entries similar to the following).**

```

06/04/01 07:57:48: Thread ID : 65536 : User Name      :
EcDsStFtpServer | Thread 65536
06/04/01 07:57:48: Thread ID : 65536 : Database Name   :
stmgtdbl_DEV07 | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : Server Name     :
f2acg01_srvr | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : DsShTSStorage: creating
the MutexVec for this thread

```

```

06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30141 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING:
DsEstUnknownError (30141) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30143 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING:
DsEstUnknownError (30143) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30139 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING:
DsEstUnknownError (30139) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30142 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING:
DsEstUnknownError (30142) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30148 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING:
DsEstUnknownError (30148) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30144 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING:
DsEstUnknownError (30144) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30145 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING:
DsEstUnknownError (30145) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30147 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING:
DsEstUnknownError (30147) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30146 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING:
DsEstUnknownError (30146) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30211 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING:
DsEstUnknownError (30211) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30140 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING:
DsEstUnknownError (30140) | Thread 65536

```

- Get server configuration parameters from the database (log entries similar to the following).

```

06/04/01 07:57:49: Thread ID : 65536 : BaseReal::Ctor: Server
Name is - EcDsStFtpServerNONE | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : User Name      :
EcDsStFtpServer | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : Database Name      :
stmgtdbl1_DEV07 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : Server Name      :
f2acg01_srvr | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : myTransactionList[0]: use
stmgtdbl1_DEV07 | Thread 65536

```



```

06/04/01 07:57:50: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: use stmgtdb1_DEV07 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : myTransactionList[1]:
exec DsStCPSelectByName "EcDsStFtpServerNONE" | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: exec DsStCPSelectByName "EcDsStFtpServerNONE" | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 :
DBIF:Fetched:[8.000000][EcDsStFtpServerNONE][1][10][FTP][][0][FTP
A][NONE][4194304] | Thread 65536

```

- **Spawn receptionist thread and register server in the database (log entries similar to the following).**

```

06/04/01 07:57:50: Thread ID : 65536 :
DsStReceptionist:BindSocketGetInfo: Port assigned is 13441 |
Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : myTransactionList[0]: use
stmgtdb1_DEV07 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: use stmgtdb1_DEV07 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : myTransactionList[1]:
exec DsStCPRegisterServer 8, 13441, "f2acg01" | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: exec DsStCPRegisterServer 8, 13441, "f2acg01" | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : Ftp:Ctor:
EcDsStFtpServerNONE | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : myTransactionList[0]: use
stmgtdb1_DEV07 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: use stmgtdb1_DEV07 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : myTransactionList[1]:
exec DsStCPSelectById 8 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: exec DsStCPSelectById 8 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 :
DBIF:Fetched:[8.000000][EcDsStFtpServerNONE][1][10][FTP][f2acg01]
[13441][FTPA][NONE][4194304] | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : Ftp:Ctor: Leaving |
Thread 65536

```

- **Spawn service threads (log entries similar to the following).**

```

06/04/01 07:57:50: Thread ID : 65536 : Ftp:Startup: temperature
= cold | Thread 65536
06/04/01 07:57:50: Performing startup processing | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : Spawning service threads
| Thread 65536
06/04/01 07:57:50: Thread ID : 65536 :
BR:GetThreadPoolConfiguration | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : myTransactionList[0]: use
stmgtdb1_DEV07 | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: use stmgtdb1_DEV07 | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : myTransactionList[1]:
exec DsStSTCSelectForServer 8, "ThreadPool" | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: exec DsStSTCSelectForServer 8, "ThreadPool" | Thread 65536

```

```

06/04/01 07:57:51: Thread ID : 65536 : DBIF:Fetch:[ThreadPool
][10][0][0][0][10] | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : 3_2709893_0757-
1125858625_155062001_f2acg01:FTPA: BR:GetThreadPoolConfiguration
Return
ing | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : Ftp: Spawning a service
thread | Thread 65536
06/04/01 07:57:51: Starting a new service thread | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : Ftp: Spawning a service
thread | Thread 65536
06/04/01 07:57:51: 06/04/01 07:57:51: Thread ID : 65554 : Waiting
for work | Thread 65554
06/04/01 07:57:51: Thread ID : 65554 : DsShTSStorage: creating
the MutexVec for this thread
06/04/01 07:57:51: Thread ID : 65554 : Waking up manager thread
| Thread 65554
.
.
.
06/04/01 07:57:51: Starting a new service thread | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : Ftp: Spawning a service
thread | Thread 65536
06/04/01 07:57:51: Starting a new service thread | Thread 65536
06/04/01 07:57:51: Thread ID : 65559 : Waiting for work | Thread
65559
06/04/01 07:57:51: Thread ID : 65560 : Waiting for work | Thread
65560
06/04/01 07:57:51: Thread ID : 65561 : Waiting for work | Thread
65561

```

- Process Restart Notification for server restart ("Ready to accept requests") (log entries similar to the following).

```

06/04/01 07:57:51: Thread ID : 65536 : myTransactionList[0]: use
stmgtddl_DEV07 | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: use stmgtddl_DEV07 | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : myTransactionList[1]:
BEGIN TRANSACTION OUTER_278888352 | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: BEGIN TRANSACTION OUTER_278888352 | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : myTransactionList[2]:
exec DsStGRRestartNotification "10_2709893_0757-1125858625_15506
2001_f2acg01:FTPA:Server restart", "EcDsStFtpServerNONE", "cold"
| Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: exec DsStGRRestartNotification "10_2709893_0757-1125858625
_155062001_f2acg01:FTPA:Server restart", "EcDsStFtpServerNONE",
"cold" | Thread 65536
06/04/01 07:57:52: Thread ID : 65536 : DBIF:Fetch:[] | Thread
65536
06/04/01 07:57:52: Thread ID : 65536 : DBIF:Fetch:[8.000000] |
Thread 65536
06/04/01 07:57:52: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: COMMIT TRANSACTION OUTER_278888352 | Thread 65536
06/04/01 07:57:52: Thread ID : 0 : No servers to awaken -- get
status | Thread 0

```

```
06/04/01 07:57:52: Thread ID : 65536 : Spawning manager thread |
Thread 65536
06/04/01 07:57:52: Ready to accept requests | Thread 65564
```

- Check queue for requests ("Waiting for an event" means there is nothing else in the queue.) (log entries similar to the following).

```
06/04/01 07:57:52: Thread ID : 65564 :
BR:ProcessCancelledRequests | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : DsShTSStorage: creating
the MutexVec for this thread
06/04/01 07:57:52: Thread ID : 65564 : myTransactionList[0]: use
stmgtdbl_DEV07 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : DBIF:Execute: Ultimate
SQL: use stmgtdbl_DEV07 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : myTransactionList[1]:
exec DsStGRSelectCancelled 8 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : DBIF:Execute: Ultimate
SQL: exec DsStGRSelectCancelled 8 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 :
BR:ProcessCancelledRequests: Nothing cancelled | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 :
BR:ProcessCancelledRequests Returning | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : Ftp: Getting next request
| Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : myTransactionList[0]: use
stmgtdbl_DEV07 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : DBIF:Execute: Ultimate
SQL: use stmgtdbl_DEV07 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : myTransactionList[1]:
exec DsStFRGetNextRequest 8 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : DBIF:Execute: Ultimate
SQL: exec DsStFRGetNextRequest 8 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : GetNextRequest: No
requests available | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : Waiting for an event |
Thread 65564
```

- The log file for the server from which the call originated may indicate a problem completing a connection. The log should indicate successful awakening of a remote host, with entries similar to the following:

```
06/04/01 07:57:52: Thread ID : 65536 :
DsStPatron:AwakenRemoteServer: Hostname - f2acg01 | Thread 65536
06/04/01 07:57:52: Thread ID : 65536 :
DsStPatron:AwakenRemoteServer: Port Number - 13441 | Thread 65536
06/04/01 07:57:52: Thread ID : 65536 : Patron: Creating new
entry for EcDsStFtpServerNONE | Thread 65536
06/04/01 07:57:52: Thread ID : 65536 : Trying gethostbyname_r()
0 of 5 attempts | Thread 65536
06/04/01 07:57:52: Thread ID : 65536 : Waking up
EcDsStFtpServerNONE | Thread 65536
```

and should indicate completion of a connection to the called server, with entries similar to the following:

```
06/04/01 07:57:52: Thread ID : 65553 :
DsStReceptionist:WaitForConnections: A connection has been
accepted | Thread 65553
```

```

06/04/01 07:57:52: Thread ID : 65564 :
BR:ProcessCancelledRequests | Thread 65564
06/04/01 07:57:52:
06/04/01 07:57:52: Thread ID : 65553 : Waking up manager thread
| Thread 65553
06/04/01 07:57:52: Thread ID : 65564 : : 06/04/01 07:57:52:
read ID : 7:57:52: DsShTSSStorageDsShTSSStorage: creating the
MutexVec for this thread: creating the MutexVec for this
thread665553: 53 : DsShTSSStorage: creating the MutexVec for this
thread
06/04/01 07:57:52: Thread ID : 65564 : myTransactionList[0]: use
stmgtdbl_DEV07 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : DBIF:Execute: Ultimate
SQL: use stmgtdbl_DEV07 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : myTransactionList[1]:
exec DsStGRSelectCancelled 8 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : DBIF:Execute: Ultimate
SQL: exec DsStGRSelectCancelled 8 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 :
BR:ProcessCancelledRequests: Nothing cancelled | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 :
BR:ProcessCancelledRequests Returning | Thread 65564.

```

- This procedure is applicable for reviewing logs for different types of errors and events on ECS servers.

5 Exit the log file (e.g., from **pg**, type **q** and then press the **Return/Enter** key).

Table 17.7-8. Checking Server Log Files

Step	What to Do	Action to Take
1	Log in at host for server and log(s) to be examined	enter text; press Return/Enter
2	cd /usr/ecs/<MODE>/CUSTOM/logs	enter text; press Return/Enter
3	pg (or other viewing command) filename	enter text; press Return/Enter
4	Review the entries in the file	read text
5	Exit the log file (e.g., q for exit from pg)	enter text; press Return/Enter

17.7.2.2 A Special Case: Checking the Request Manager Server Debug Log

The Request Manager server in the Storage Management computer software configuration item of the Data Server Subsystem processes requests from external clients (processes outside of Storage Management). Requests between Storage Management servers are passed directly from one server to another.

- Requests that require one of the Storage Management servers to perform processing are checkpointed (except requests that can be serviced solely through SQL).
 - Checkpointing involves recording the request's state (e.g., "checkpointed," "failed," "completed") in the database to assist in error recovery.
- Requests that can be serviced solely through SQL are considered "trivial" requests.

- Trivial requests are not checkpointed.
- Examples include attaching to a staging disk, getting capacity, and getting block size.
- Trivial requests submitted from outside Storage Management are serviced by the Request Manager server.
- Trivial requests originating within Storage Management are passed directly from the client to the database server.

The Request Manager server (like other Storage Management servers) can manage several concurrent activities. This is accomplished through the use of threads. There are several different kinds of threads:

- Manager thread.
 - One per Storage Management server.
 - Responsible for dequeuing requests and assigning them to service threads.
 - Checks for cancelled requests.
- Service thread.
 - Multiple threads per Storage Management server.
 - Responsible for the actual servicing of requests.
 - Logs all progress including all changes of request state.
 - Notifies submitter when request has been completed.
- Receptionist thread.
 - One per Storage Management server.
 - Registers the server as "up" in the database.
 - Sits on a socket, waiting for connections from other Storage Management servers.
 - Unregisters the server at shutdown.
- Inbound RPC thread.
 - Spawned by a request from a Storage Management client.
 - Hands off the request to the manager thread and waits for completion of the request.
- Housekeeper thread.

- Watches for completed requests which haven't previously been seen and processed.

Information concerning Request Manager server processing of requests (identified by thread) is recorded in the Request Manager server debug log (assuming some level of debug log recording is specified in the Registry database).

Trivial requests typically involve the following types of activities:

- Inbound RPC thread appears with a request.
- Manager thread dequeues the request and assigns it to a service thread.
- Service thread recognizes the thread as "trivial."
 - A "No checkpointing required -- going straight to responded" message is recorded in the Request Manager server debug log.
- Service thread executes the database transaction for results.
 - When the request is completed, a "Done servicing" message is recorded in the Request Manager server debug log.
 - If the request fails, an "Unable to service" message is recorded in the Request Manager server debug log.
- Service thread hands the results to the inbound RPC thread.
 - A "Notifying the client" message is recorded in the Request Manager server debug log.
- Inbound RPC thread silently returns to the client with the results.

Non-trivial requests are forwarded to the appropriate Storage Management server (e.g., EcDsStFtpServer, EcDsStStagingDiskServer, EcDsStArchiveServer) for processing.

- Some of the same types of entries are made in the Request Manager server debug log for non-trivial requests as for trivial requests.
 - For example:
 - "Waking up service thread" (Request Manager is preparing to process the request).
 - "Done servicing" (request processing has been completed).
 - "Unable to service" (the request has failed).
- Although some trivial requests include "token" statements, tokens are characteristic of non-trivial requests.
 - A token includes request information that varies with the type of operation to be performed.

- For example, a token for an ftp request might include the following types of data:
 - Stored procedure (e.g., DsStFRInsert) [other types of stored procedures include DsStSDRInsert and DsStGRMapLogicalArchiveId].
 - R P C I D (e . g . , R P C I d = 1 8 2 1 _ 5 3 5 _ 1 1 0 9 - 1124464729_171062001_x0ins01.xdc.ecs.nasa.gov:SBSVSDSV1DSDD 1DSDD4:).
 - Username.
 - Encrypted password.
 - Host.
 - Source path.
 - Destination path.
 - External request ID.
 - Server name (e.g., EcDsStFtpServerNONE) [other types of operations might involve the EcDsStStagingDiskServerDRP1 for example].
 - Type of operation (e.g., FtpPush) [other types of operations include ArRetrieve, SDAllocateDisk, SDLINKFile].
 - Submitter (e.g., DSDD) [other types of operations might involve SDSV].
 - Priority.
- The server to which the request was sent is identified by name (ServerName).
- Transaction ID is embedded in the RPC ID (the portion before the first colon in the RPC ID).

A "transaction" may involve multiple operations on a host or several hosts. Consequently, multiple threads may be used on each relevant host.

Table 17.7- 9 presents the general steps required for checking the Storage Management Request Manager server debug log file. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in to the Distribution Server host (e.g., e0dis02, g0dis02, l0dis02, n0dis02).
- 2 To change to the logs directory, type **cd /usr/ecs/<MODE>/CUSTOM/logs** then press the **Return/Enter** key.
 - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/logs**.
- 3 Type **pg filename** then press the **Return/Enter** key.
 - *filename* refers to the appropriate Request Manager debug log.

- For example: **pg EcDsStRequestManagerServerDebug.log**
- The content of the first page of the specified file is displayed.
- Although this procedure has been written for the **pg** command, any UNIX editor or visualizing command (e.g., **vi**, **view**, **more**) can be used to review the log file.

4 At the **:** prompt type **/date time** then press the **Return/Enter** key.

- **date time** refers to the approximate date and time of the problem.
 - For example:
/06/18/01 12:17:31
- The file is searched for the specified text.
 - If the specified text is in the log file, the following type of response is displayed.

```
...skipping forward
06/18/01 12:17:31: Thread ID : 105 : DsShTSStorage: creating the
MutexVec for this thread
[...]
```
 - If the specified text is not in the log file, the following type of response is displayed.

```
Pattern not found:
```
 - If the specified text is not in the log file, verify the following aspects of Steps 3 and 4:
 - Date and time were entered correctly (Step 4).
 - Proper file was opened (Step 3).

5 At the **:** prompt type **/Unable to service** then press the **Return/Enter** key.

- **pg** searches the file for the specified text.
 - If the specified text is in the log file, the following type of response is displayed.

```
...skipping forward
2:IngestRQ409GR1 Unable to service | Thread 52
[...]
```
 - If the specified text is not in the log file, the following type of response is displayed.

```
Pattern not found:
```


- If the specified text is in the file, go to Step 7.
 - If the specified text is not in the file, go to Step 6.
- 6** Examine the contents of the log file to determine which thread is associated with the problem being investigated.
- The following **pg** commands (at the **:** prompt) are useful:
 - ***n*** then **Return/Enter** (go to Page *n*).
 - **Return/Enter** or **+1** then **Return/Enter** (go down to the next page).
 - **-1** then **Return/Enter** (go back to the preceding page).
 - **+*n*** then **Return/Enter** (go down *n* number of pages).
 - **-*n*** then **Return/Enter** (go back *n* number of pages).
 - **+*nl*** then **Return/Enter** (go down *n* number of lines).
 - **-*nl*** then **Return/Enter** (go back *n* number of lines).
 - **q** then **Return/Enter** (exit from **pg**).
- 7** At the **:** prompt type the appropriate text (depending on the direction of the desired search) then press the **Return/Enter** key:
- To search back toward the beginning of the file, type **^Waking up service thread *n*^** and then press **Return/Enter**.
- To search toward the end of the file, type **/Waking up service thread *n*** and then press **Return/Enter**.
- For example:
 - ^Waking up service thread 52^**
 - The file is searched back toward the beginning of the file for the specified text.
 - If the specified text is in the log file, the following type of response is displayed.


```
...skipping backward
06/18/01 12:17:31: Thread ID : 102 : Waking up service thread 52
| Thread 102
[...]
```
 - If the specified text is not in the log file, the following type of response is displayed.


```
Pattern not found:
```

- The entries "Waking up service thread *n*" and "Unable to service | Thread *n*" bracket the thread servicing in which an error occurred.

NOTE: Thread IDs are reused frequently. There are likely to be many processes with the same thread ID in any particular log file. It is important to follow the correct instance of the thread.

NOTE: It is likely that the Request Manager would try again to process a failed request. Subsequent request processing may use the same thread ID or a different thread ID. However, it would involve the same transaction ID.

- A "No checkpointing required -- going straight to responded" entry associated with the thread ID indicates that the request is "trivial."

8 At the : prompt type **/SEARCHING** then press **Return/Enter**.

- The file is searched for the specified text.
 - If the specified text is in the log file, the following type of response is displayed.

```
...skipping forward
06/18/01 12:17:31: Thread ID : 52 : SEARCHING FOR: 30148 (Found)
| Thread 52
06/18/01 12:17:31: Thread ID : 52 : SEARCHING FOR: 30148 (Found)
| Thread 52
06/18/01 12:17:31: Thread ID : 52 :
DsStStoredProcedures::Execute - ERROR: Could not execute stored
procedure | Thread 52
06/18/01 12:17:31: Thread ID : 52 : Error encountered in stored
procedure | Thread 52
06/18/01 12:17:31: Thread ID : 52 : DBIF:Execute: Ultimate SQL:
ROLLBACK TRANSACTION OUTER_7077776 | Thread 52
06/18/01 12:17:32: Thread ID : 52 : 1_4501810_1217-
1124633447_169062001_p0icg01.pvc.ecs.nasa.gov:IPOBIPOB1INRM1IGSA1
5:IngestRQ409GR1 Done servicing | Thread 52
06/18/01 12:17:32: Thread ID : 52 : 1_4501810_1217-
1124633447_169062001_p0icg01.pvc.ecs.nasa.gov:IPOBIPOB1INRM1IGSA1
5:IngestRQ409GR1 Unable to service | Thread 52
06/18/01 12:17:32: Thread ID : 52 : 1_4501810_1217-
1124633447_169062001_p0icg01.pvc.ecs.nasa.gov:IPOBIPOB1INRM1IGSA1
5:IngestRQ409GR1 Marked as unassigned | Thread 52
06/18/01 12:17:32: Thread ID : 52 : 1_4501810_1217-
1124633447_169062001_p0icg01.pvc.ecs.nasa.gov:IPOBIPOB1INRM1IGSA1
5:IngestRQ409GR1 Notifying the client | Thread 52
06/18/01 12:17:32: Thread ID : 52 : Waiting for work | Thread 52
06/18/01 12:17:32: Thread ID : 52 : Waking up manager thread |
Thread 52
[...]
```

- In the preceding example the expression **SEARCHING** is associated with Thread ID 52.

- The context of the **SEARCHING** statement indicates the type and source of the problem; in this case there appears to be a problem executing a stored procedure.

- If the specified text is not in the log file, the following type of response is displayed.

Pattern not found:

9 If the expression **SEARCHING** is not associated with the specified thread in the lines displayed, repeat Step 8.

10 If necessary, at the : prompt type **-2l** [lower-case letter l] then press the **Return/Enter** key.

- **pg** simulates scrolling the screen backward two lines (or any other number of lines that is typed at the prompt).
 - The file is redisplayed to include the two lines that preceded the page previously displayed.
 - For example:

```
...skipping backward
06/18/01 12:17:31: Thread ID : 52 : DBIF:Execute: Ultimate SQL:
exec DsStSDAttachDisk
"/usr/ecs/TS2/CUSTOM/pdps/x0spg01/data/DpPrRm/x0spg01_disk",
"SDSV", 0 | Thread 52
06/18/01 12:17:31: Thread ID : 52 : SEARCHING FOR: 30148 (Found)
| Thread 52
06/18/01 12:17:31: Thread ID : 52 : SEARCHING FOR: 30148 (Found)
| Thread 52
06/18/01 12:17:31: Thread ID : 52 :
DsStStoredProcedures::Execute - ERROR: Could not execute stored
procedure | Thread 52
06/18/01 12:17:31: Thread ID : 52 : Error encountered in stored
procedure | Thread 52
[...]
```

- The additional lines preceding "SEARCHING FOR" in the example indicate that the stored procedure in which the error was encountered is DsStSDAttachDisk.

11 At the : prompt type **q** then press the **Return/Enter** key.

- **pg** exits from the Request Manager server debug log file.

- 12 If the request is a trivial request, go to Step 22.
- 13 If the request is a non-trivial request, open a separate UNIX window.
 - The results of related operations on the server involved in performing copy or ftp functions for the transaction are going to be checked in a separate UNIX window.
- 14 In the new UNIX window log in to the appropriate server host (e.g., e0drg11, g0drg01, l0drg01, n0drg01) for the server involved in performing copy or ftp functions for the transaction.
- 15 At the shell prompt type **grep 'TransactionId' filename | grep 'LogProgress'** then press the **Return/Enter** key.
 - For example:


```
grep 'af610628-' EcDsStArchiveServerDebug.log | grep 'LogProgress'
```
 - *filename* refers to the name of the log file for the process involved in performing copy or ftp functions for the transaction.
 - *TransactionId* refers to the Transaction ID associated with the applicable request.
 - In this example af610628-1dd1-11b2-a047-af3a589fd88e is the relevant Transaction ID.
 - However, usually it is not necessary to use the entire Transaction ID in the command; a representative sample (e.g., af610628- from the example) should be sufficient.
 - References to other Transaction IDs and entries that do not contain the string "LogProgress" are filtered out so references to the specified Transaction ID that contain the string "LogProgress" are the only log entries displayed.
 - The string "LogProgress" is a filter for references to stored procedure DsStGRLogProgress.
 - Progress is logged for copy and ftp input/output at each block.
 - The following type of response is displayed:

```
06/26/01 12:46:00: Thread ID : 65674 : myTransactionList[1]:
exec DsStGRLogProgress "af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2
SC:MOD03.001:55732", 0, 1, "files" | Thread 65674
06/26/01 12:46:00: Thread ID : 65674 : DBIF:Execute: Ultimate
SQL: exec DsStGRLogProgress "af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2
SC:MOD03.001:55732", 0, 1, "files" | Thread 65674
06/26/01 12:46:43: Thread ID : 65674 : : 06/26/01 12:46:43:
read ID : 2:46:43: myTransactionmyTransactionList[1]: exec
DsStGRLogProgress "af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2
```

```
SC:MOD03.001:55732", 60, 60, "MB"List[1]: exec DsStGRLogProgress
"af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2
SC:MOD03.001:55732", 60, 60, "MB"65714read 65674 : 74
06/26/01 12:46:43: Thread ID : 65674 : DBIF:Execute: Ultimate
SQL: exec DsStGRLogProgress "af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2
SC:MOD03.001:55732", 60, 60, "MB"0DBIF:Execute: Ultimate SQL:
exec DsStGRLogProgress "af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2
SC:MOD03.001:55732", 60, 60, "MB"06/26/01 12:46:43: 6/26/01
12:46:43: | Thread : 65714read 65674 : 74
```

- If no progress is indicated, go to Step 22.

16 Click in the UNIX window for the Distribution Server host.

17 Type **grep 'TransactionId' filename | grep 'Done servicing'** then press **Return/Enter**.

- *filename* refers to the appropriate Request Manager debug log.
- For example:

```
grep 'af610628-' EcDsStRequestManagerServerDebug.log | grep 'Done servicing'
```

- If the operation has been completed, the following type of response is displayed:

```
06/26/01 12:46:00: Thread ID : 52 : af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2
SC:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:44: Thread ID : 52 : af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2
SC:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:45: Thread ID : 52 : af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD2DSDD1DSDD3:MoPGE02#sy14182000TS2S
C:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:47: Thread ID : 52 : af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD2DSDD1DSDD3:MoPGE02#sy14182000TS2S
C:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:47: Thread ID : 52 : af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD2DSDD1DSDD7:MoPGE02#sy14182000TS2S
C:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:50: Thread ID : 52 : af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD2DSDD1DSDD7:MoPGE02#sy14182000TS2S
C:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:51: Thread ID : 52 : af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD4:MoPGE02#sy14182000TS2SC:MOD03.00
1:55732 Done servicing | Thread 52
06/26/01 12:46:56: Thread ID : 52 : af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD4:MoPGE02#sy14182000TS2SC:MOD03.00
1:55732 Done servicing | Thread 52
```

```
06/26/01 12:46:56: Thread ID : 52 : af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD8:MoPGE02#sy14182000TS2SC:MOD03.00
1:55732 Done servicing | Thread 52
06/26/01 12:46:59: Thread ID : 52 : af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD8:MoPGE02#sy14182000TS2SC:MOD03.00
1:55732 Done servicing | Thread 52
```

- The statement "Done servicing" shows that the operation has been completed; however, it provides no indication as to whether the operation succeeded or failed.
- If "Done servicing" is followed by "Unable to service," (as described in Step 19) the operation failed.
- If the operation has not been completed, no file entries are displayed (the UNIX prompt is displayed).
 - It may just be slow to complete.
- If the operation has been completed, go to Step 19.
- If the operation has not been completed, go to Step 20.

18 Type **grep 'TransactionId' filename | grep 'Unable to service'** then press the **Return/Enter** key.

- *filename* refers to the appropriate Request Manager debug log.
- For example:
grep '2a7d4168-' EcDsStRequestManagerServerDebug.log | grep 'Unable to service'
- If the request has failed, the following type of response is displayed:

```
06/26/01 12:56:22: Thread ID : 52 : 2a7d4168-1dd2-11b2-8c52-
99d0f708dce5:PDPSSDSV1:MoPGE02#sy14182000TS2MOD02OBC Unable to
service | Thread 52
06/26/01 12:56:22: Thread ID : 52 : 2a7d4168-1dd2-11b2-8c52-
99d0f708dce5:PDPSSDSV4:MoPGE02#sy14182000TS2MOD02OBC Unable to
service | Thread 52
```

- If the operation has failed, return to Step 7.
- If the operation has not failed, no file entries are displayed (the UNIX prompt is displayed).

19 If the operation has not failed, at the shell prompt type **tail -f filename | grep 'TransactionId'** and then press the **Return/Enter** key.

- *filename* refers to the appropriate Request Manager debug log.
- *TransactionId* refers to the Transaction ID associated with the applicable request.

- For example:
tail -f EcDsStRequestManagerServerDebug.log | grep 'af610628-'
 - If new entries are being posted to the log, the operation has not finished yet.
 - If the same entries continue to be repeated over and over, it may be necessary to restart the server.
 - If it is necessary to exit from a tailed log, type **ctrl-c** (while holding down the **Control Key**, press **c**).
- 20** If the operation has not finished yet, monitor the tailed log for awhile.
- If the operation does not seem to finish (i.e., if entries continue to be made to the tailed log) after a reasonable period of time (e.g., 30 minutes), it may be necessary to restart the Request Manager server.
 - If it is necessary to exit from a tailed log, type **ctrl-c** (while holding down the **Control Key**, press **c**).
- 21** If problems were detected in the Request Manager server debug log and/or the log file for the process involved in performing copy or ftp functions for the transaction, it may be necessary to restart the server(s) performing those functions.
- If server restart does not resolve the problem, it is appropriate to notify the Help Desk and prepare a Trouble Ticket.
- 22** If no problems were detected in the Request Manager server debug log or the log file for the process involved in performing copy or ftp functions for the transaction, check the Science Data Server log files; use Procedure 17.7.2.1 for **Checking Server Log Files**.

Table 17.7- 9. A Special Case: Checking the Request Manager Server Debug Log

Step	What to Do	Action to Take
1	Log in at host for Distribution Server	enter text; press Return/Enter
2	cd /usr/ecs/<MODE>/CUSTOM/logs	enter text; press Return/Enter
3	pg (or other viewing command) EcDsStRequestManagerServerDebug.Log (or <i>filename</i> of other Request Manager debug log)	enter text; press Return/Enter
4	/date time (of problem)	enter text; press Return/Enter
5	/Unable to service	enter text; press Return/Enter
6	Identify thread ID associated with problem	read text
7	^Waking up service thread n^ or /Waking up service thread n	enter text; press Return/Enter
8	/SEARCHING	enter text; press Return/Enter
9	As necessary, repeat Step 8	
10	As necessary, -2l	enter text; press Return/Enter
11	q (to exit pg)	enter text; press Return/Enter
12	If request is trivial, go to Step 22	
13	If request is non-trivial, open a separate UNIX window	enter text; press Return/Enter
14	(In new window), log in at host for server for transaction	enter text; press Return/Enter
15	grep 'TransactionId' filename grep 'LogProgress'	enter text; press Return/Enter
16	Go to window for Distribution Server host	single-click
17	(In logs directory) grep 'TransactionId' filename grep 'Done servicing'	enter text; press Return/Enter
18	grep 'TransactionId' filename grep 'Unable to service'	enter text; press Return/Enter
19	If operation has not failed, tail -f filename grep 'TransactionId'	enter text; press Return/Enter
20	If operation not finished, monitor tailed log	read text
21	If problem(s) detected, restart associated server	
22	If no problem detected, check Science Data Server logs	Use Procedure 17.7.2.1

17.7.2.3 Checking the **el_ETAC** Log

Each day a current **el_ETAC_00** log on the FSMS host records interactions between AMASS and ACSLS. This log can provide helpful information in troubleshooting problems manifested in those interactions. Table 17.7-10 presents the steps required for recovering from a failure to store data. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in as **amass** at the FSMS host.
- 2 Type **cd /usr/amass/emasslogs/etac** and then press the **Return/Enter** key.
 - The working directory is changed to **/usr/amass/emasslogs/etac**.
- 3 Use the current el_ETAC log to investigate possible problems in communication between AMASS and ACSLS. To view the current el_ETAC log, type **pg el_ETAC_00** and then press the **Return/Enter** key.
 - The first page of the log file is displayed; additional sequential pages can be displayed by pressing the **Return/Enter** key at the **:** prompt.
 - Although this procedure has been written for the **pg** command, any UNIX editor or visualizing command (e.g., **vi**, **view**, **more**, **tail**) can be used to review the log file.
 - The log contains entries related to activities and communications associated with actions by AMASS to direct ACSLS robotic activities; the entries should appear in format similar to the following sample:

```
Sep 24 09:49:42 p0drg01 amass LIBSCHED3[7215638]:  
E7003(16)<00000>:xdiStk2749: STK Response received; Status: 0  
  
Sep 24 09:49:42 p0drg01 amass LIBSCHED3[7215638]:  
E7003(16)<00000>:xdiStk2797: ACSLS ACK response received  
  
Sep 24 09:49:42 p0drg01 amass LIBSCHED3[7215638]:  
E7003(16)<00000>:xdiStk2742: Waiting for ACSLS response  
  
Sep 24 09:49:51 p0drg01 amass LIBSCHED3[7215638]:  
E7003(16)<00000>:xdiStk2749: STK Response received; Status: 0  
  
Sep 24 09:49:51 p0drg01 amass LIBSCHED3[7215638]:  
E7003(16)<00000>:xdiStk2873: ACSLS final response received  
  
Sep 24 09:49:51 p0drg01 amass LIBSCHED3[7215638]:  
E7003(16)<00000>:xdiStk2876: 1 network packets transfered  
  
Sep 24 10:18:52 p0drg01 amass LIBSCHED1[7215870]:  
E7003(16)<00000>:xdiArch39: Archive index : 0  
  
Sep 24 10:18:52 p0drg01 amass LIBSCHED1[7215870]:  
E7003(16)<00000>:xdiStk486: Media Id = P10011, Drive index = 0  
  
Sep 24 10:18:52 p0drg01 amass LIBSCHED1[7215870]:  
E7003(16)<00000>:xdiStk533: Sending a mount command  
  
Sep 24 10:19:32 p0drg01 amass LIBSCHED1[7215870]:  
E7003(16)<00000>:xdiArch39: Archive index : 0  
  
Sep 24 10:19:32 p0drg01 amass LIBSCHED1[7215870]:  
E7003(16)<00000>:xdiStk486: Media Id = P20676, Drive index = 1  
  
Sep 24 10:19:32 p0drg01 amass LIBSCHED1[7215870]:  
E7003(16)<00000>:xdiStk533: Sending a mount command
```

```

Sep 24 10:34:56 p0drg01 amass LIBSCHED1[7215870]:
E1043(7)<00000>:libsched3165: Idle Eject timer expired on volume
188 in drive 2.

Sep 24 10:35:07 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiArch39: Archive index : 0

Sep 24 10:35:07 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiStk686: Media Id = P20676, Drive index = 1

Sep 24 10:35:07 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiStk719: Sending a dismount command

Sep 24 10:35:07 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiStk2742: Waiting for ACSLS response

Sep 24 10:35:07 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiStk2749: STK Response received; Status: 0

Sep 24 10:35:07 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiStk2758: Error unexpected sequence number: 101
-expected sequence number: 109

```

- Examine the sections of the log with entries near the time of any problem being investigated, looking for messages that indicate whether there was successful communication between AMASS and ACSLS regarding mounting of a tape and transfer of information. It may be useful to search the log for occurrences of the word **fail** (while viewing the log with **pg**, **view**, **vi**, or other viewing/editing tool, type **/fail** and press the **Return/Enter** key).
- If the log indicates problems in communication between AMASS and ACSLS, it may be useful to use the **quedisplay** command to obtain the AMASS view of the queue and the **medialist** command to obtain the robot view. If these commands show discrepancies indicating a lack of synchrony between AMASS and ACSLS, it may be possible to re-establish that synchrony using the **mediamove** command (refer to Procedure 17.7.1.6 **Using *mediamove* to Establish Synchrony Between *quedisplay* and *medialist***).
- **Note:** The message "Error unexpected sequence number: 101 -expected sequence number: 109" is an artifact likely to be removed in releases of AMASS subsequent to Version 5.0.0 Revision 17 and does not reflect a real error.

Table 17.7-10. Checking the el_ETAC Log

Step	What to Do	Action to Take
1	Log in at FSMS host	enter text; press Return/Enter
2	cd /usr/amass/emasslogs/etac	enter text; press Return/Enter
3	pg el_ETAC_00 (or vi , view , tail , or other viewing tool)	enter text; press Return/Enter

17.7.2.4 Handling a Data Insertion Failure

Successful data insertion requires interactions among numerous servers, and the interactions are reflected in entries in the debug logs for those servers. Detection and initial isolation of a problem that prevents successful insertion may require tracing events across multiple log files on different hosts. Table 17.7-11 presents the steps required for recovering from a failure to store data. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 4 At the host for SDSRV (e.g., e0acs05, g0acs03, l0acs03, n0acs04), review the debug log **EcDsScienceDataServerDebug.log** (use Procedure 17.7.2.1 **Checking Server Log Files**).
 - Examine the section of the log with entries near the time of the problem, looking for error messages that indicate communication failure.
 - If the log file entries indicate a communication problem, note the server(s) with which there is impairment or disruption of communication.
- 5 At the host for Archive Server (e.g., e0drg11, g0drg01, l0drg01, n0drg01), review the debug log **EcDsStArchiveServerDebug.log** (use Procedure 17.7.2.1 **Checking Server Log Files**).
 - Examine the section of the log with entries near the time of the problem, looking for error messages that indicate communication failure.
 - If the log file entries indicate a communication problem, note the server(s) with which there is impairment or disruption of communication.
- 6 If Step 1 and/or Step 2 resulted in detection of a problem in the interaction of SDSRV and/or Archive Server with other servers, at the host(s) for those servers, review the server debug log(s). These logs may include:
 - EcDsStStagingDiskServerDebug.log** (on Archive Server host).
 - EcDsStCacheManagerServerDebug.log** (on Archive Server host).
 - EcDsStRequestManagerServerDebug.log** (e.g., on e0dis02, g0dis02, l0dis02, n0dis02; use Procedure 17.7.2.2).
 - EcIoAdServerDebug.log** (e.g., on e0ins02, g0ins02, l0ins02, n0ins02).
 - EcSbSubServerDebug.log** (e.g., on e0ins01, g0ins01, l0ins01, n0ins01).
 - If there is evidence of requests not succeeding or other communication failure, it may be necessary to have System Administrators or Engineering Support personnel resolve the problem (e.g., restart affected servers, execute **EcCsIdPingServers**, ensure that the **Name Server** is up in the mode being used and that its debug log reflects appropriate look-up activity by the application servers, mount points are intact, and database access is not impaired).

Note: The next three steps address running the Check Archive script, **EcDsCheckArchive**. To run this script, it is necessary to enter eight database-specific parameters when prompted during the running of the script: STMGT SQL server name, STMGT database name, STMGT SQL server userID, STMGT SQL server database password, SDSRV SQL server name, SDSRV database name, SDSRV SQL server userID, and SDSRV database password. To facilitate the smooth execution of the script, the parameters may be set as environmental variables instead. The parameters are not readily available to most operators; therefore, you will need to obtain them from the Database Administrator or have the Database Administrator run the script for you, using steps 4 through 6.

4 On the host for the Archive Server, type **cd /usr/ecs/<MODE>/CUSTOM /utilities** and then press the **Return/Enter** key.

- The prompt reflects the directory change to **/usr/ecs/<MODE>/CUSTOM/utilities**.

5 Type **EcDsCheckArchive <MODE>**.

- The Check Archive script runs; the initially displayed information should be similar to the following:

```
=====
This script is designed to validate the Inventory
against the Archive.

The user must select the menu option associated with the
Volume Group to be validated

Please press [RETURN] to continue
=====
```

6 Follow the on-screen prompts for the script, entering the necessary parameters.

- The script provides indication of any discrepancies between the presence of granules in the Archive and entries in the inventory (metadata). Note that the appearance of a discrepancy is not necessarily indication of a failure (e.g., if a granule has been deleted but the inventory database has not been cleaned up, there may be inventory entries for which there are no granules in the archive), but a problem may be indicated if a discrepancy is apparent for a granule that you just inserted. Note also that this script would not reveal a problem if you attempted to insert a granule which failed to get inserted and also had its metadata fail to be inserted into the inventory (i.e., no granule and no inventory entry = no discrepancy). Therefore, if the script reveals no discrepancies, it may still be useful to conduct a direct examination to determine if the granule has been inserted.

7 On the host for the Archive Server, type the directory change command **cd /dss_stk1/<MODE>/<data_type_directory>** and then press the **Return/Enter** key.

- The working directory is changed to **/dss_stk1/<MODE>/<data_type_directory>**.

- 8 Type **ls -al | grep "<date>"** where "<date>" is a three-letter abbreviation for the month followed by a number indicating the day (e.g., "**Apr 21**") for the granule being inserted, and then press the **Return/Enter** key.
 - If the inserted file is displayed, with date and time of entry, go to Step 9.
 - If the inserted file is not displayed, have the Ingest/Distribution Technician insert the file again. If this succeeds (i.e., the file is now listed), go to Step 9; otherwise, conduct the procedure for **Diagnosing/Investigating Write Errors** Procedure 17.7.3).
- 9 Determine if the inserted file is reflected in the Inventory Database (Database Administrator function) by logging into Sybase on the host for SDSRV and then selecting the data type for the granule being inserted.
 - If the inserted file is reflected in the Inventory Database, go to Step 10.
 - If the inserted file is not reflected in the Inventory Database, ensure that database access is not impaired (Database Administrator function).
- 10 Determine if the directory from/to which the copy is being made is visible on the machine being used; have the System Administrators or Engineering Support personnel check the mount points on the Archive host and the SDSRV host.
 - If the mount points are OK, go to Step 11.
 - If necessary, have the System Administrators or Engineering Support personnel re-establish the mount point(s).
- 11 If you inserted the file with the DSS Driver, go to Step 13. If you used Ingest to insert the file, on the Ingest host (e.g., e0icg11, g0icg01, l0icg01, n0icg01) examine the **drp**- or **icl**-mounted staging directory to determine if a staging disk was created. To do this, first type **cd /usr/ecs/<MODE>/CUSTOM/drp/<host>/data/staging/cache** (or type **cd /usr/ecs/<MODE>/CUSTOM/icl/<host>/data/StagingArea/cache**), then press the **Return/Enter** key.
 - The prompt reflects a change to the specified directory. [**Note:** Be sure that you are checking the correct mount/host. Most ingests use Ingest subsystem staging areas (i.e., **icl**), but others may not. Media ingest (e.g., from tape) typically involves staging in a **dip** area. For a polling ingest for data from EDOS, the polling directory may serve as the staging area. Some data are staged directly to working storage in the Data Server subsystem. If in doubt, consult Ingest/Archive personnel.]
- 12 Type **ls -al | more** and then press the **Return/Enter** key.
 - Any staging areas are listed in output similar to the following sample:


```
-rw-rw-r--      1 cmshared cmshared   10375 Jan 30 14:46
:SC:L70RF2.002:16015:6.HDF-EOS
-rw-rw-r--      1 cmshared cmshared   535563 Jan 30 14:46
:SC:L70RF2.002:16015:7.HDF-EOS
-rw-rw-r--      1 cmshared cmshared   154399 Jan 25 12:34
:SC:L7CPF.002:13835:1.ASCII
```

```

-rw-rw-r--      1 cmshared cmshared  154399 Jan 25 14:17
:SC:L7CPF.002:16644:1.ASCII
-rw-rw-r--      1 cmshared cmshared  154399 Jan 25 17:31
:SC:L7CPF.002:16769:1.ASCII
-rw-rw-r--      1 cmshared cmshared   67466 Jan 25 18:11
:SC:L7IGS.001:16789:1.ASCII
-rw-rw-r--      1 cmshared cmshared   43570 Jan 25 18:04
:SC:L7IGS.001:16790:1.ASCII
-rw-rw-r--      1 cmshared cmshared 499804704 Feb  6 11:49
:SC:MOD000.001:11856:1.CCSDS
-rw-rw-r--      1 cmshared cmshared 320663592 Feb  6 11:51
:SC:MOD000.001:11856:2.CCSDS
-rw-rw-r--      1 cmshared cmshared    540 Feb  6 11:51
:SC:MOD000.001:11856:3.CCSDS.

```

- If a staging area for the inserted file appears at the end of the list, go to Step 13.
 - If no staging area appears for the inserted file, it is possible that the ingest failed and that the staging area was immediately removed as part of clean-up. Check the Ingest logs (e.g., **EcInReqMgrDebug.log**, **EcInAutoDebug.log**, **EcInGranDebug.log**, or **EcInGranDebug.log**, depending on the type of Ingest) (refer to procedures for troubleshooting Ingest problems, Chapter 16) to determine if a staging disk was created. If no staging disk was created, it may be necessary to resolve a communications failure as described in Step 7.
- 13** Ensure that the Archive volume groups are set up correctly (refer to Procedure 17.3.2 **Using Storage Management GUIs to Display Archive Path Information**).
- 14** Ensure that the volume groups are on line (refer to Procedure 17.7.1.3 **Using *vollist* to Display Volume Data**).
- If the volume groups are set up correctly and their volumes are on line, and insertion still fails, it is appropriate to contact the Help Desk and prepare a trouble ticket (see Chapter 8).

Table 17.7-11. Handling a Data Insertion Failure

Step	What to Do	Action to Take
1	Review Science Data Server Debug Log	use Procedure 17.7.2.1
2	Review Archive Server Debug Log	use Procedure 17.7.2.1
3	Review debug logs for any implicated servers: a. EcDsStStagingDiskServerDebug.log b. EcDsStCacheManagerServerDebug.log c. EcDsStRequestManagerServerDebug.log d. EcIoAdServerDebug.log e. EcSbSubServerDebug.log	a. use Procedure 17.7.2.1 b. use Procedure 17.7.2.1 c. use Procedure 17.7.2.2 d. use Procedure 17.7.2.1 e. use Procedure 17.7.2.1
4	(On Archive Server host) cd /usr/ecs/<MODE>/CUSTOM/utilities	enter text; press Return/Enter
5	EcDsCheckArchive <MODE>	enter text; press Return/Enter
6	Respond to prompts to complete Check Archive Script	read text; enter text; press Return/Enter
7	(On Archive Server host) cd /dss_stk1/<MODE>/<data_type_directory>	enter text; press Return/Enter
8	ls -al grep "<date>"	enter text; press Return/Enter
9	Check for file in SDSRV Inventory database	(Database Administrator task)
10	Check mount points on Archive host and SDSRV host	(System Administrator task)
11	(For Ingested file, on Ingest host) cd /usr/ecs/<MODE>/CUSTOM/drp/<host>/data/staging/cache OR cd /usr/ecs/<MODE>/CUSTOM/icl/<host>/data/StagingArea/cache)	enter text; press Return/Enter
12	ls -al more	enter text; press Return/Enter
13	Check Archive path information to ensure volume groups are set up correctly	use Procedure 17.3.2
14	Use vollist to ensure that volume groups are online	use Procedure 17.7.1.3

17.7.2.5 Handling a Data Acquire Failure

As a first check, it is appropriate to determine if the acquire request appears in the list of System Requests on the Science Data Server GUI. If the acquire request does not appear on the Science Data Server GUI, you will need to determine where the breakdown occurred. Diagnosing an acquire failure requires detailed examination of the following system log files and directories associated with the process:

- Science Data Server log file (EcDsScienceDataServerDebug.log).
- Archive Server log file (EcDsStArchiveServerDebug.log).
- STMGIT Request Manager Server log file (EcDsStRequestManagerDebug.log)
- Staging Area.
- Presence of the relevant file.
- Staging Disk log files (EcDsStStagingDiskServerDebug.log or EcDsStCacheManagerServerDebug.log).
- Space available in the staging area.

In addition, note that a number of servers, clients, or other software running on various hosts, as reflected in Table 17.7-12, may be involved at various times in processing an acquire request. More information useful in troubleshooting may appear in related logs on these hosts.

Table 17.7-12. Hosts, Servers, Clients and Other Software Relevant to Acquires

HOST	SERVER/CLIENT/OTHER SOFTWARE
Distribution Server (e.g., e0dis02, g0dis02, l0dis02, n0dis02)	Distribution Server (EcDsDistributionServer) Request Manager Server (EcDsStRequestManagerServer)
Working Storage (e.g., e0wkg01)	Archive Server (EcDsStArchiveServer) Cache Manager Server (EcDsStCacheManagerServer) Staging Disk Server (EcDsStStagingDiskServer)
SDSRV Server (e.g., e0acs05, g0acs03, l0acs03, n0acs04)	Science Data Server (EcDsScienceDataServer) HDF EOS Server (EcDsHdfEosServer)
Access/Process Coordinators (APC) Server (e.g., e0acg11, g0acg01, l0acg02, n0acg01)	Archive Server (EcDsStArchiveServer) FTP Server (EcDsStFtpServer) Cache Manager Server (EcDsStCacheManagerServer) Staging Disk Server (EcDsStStagingDiskServer) Pull Monitor Server (EcDsStPullMonitorServer)
FSMS Server (e.g., e0drg11, g0drg01, l0drg01, n0drg01)	Archive Server (EcDsStArchiveServer) Cache Manager Server (EcDsStCacheManagerServer) Staging Disk Server (EcDsStStagingDiskServer)
Interface Server 02 (e.g., e0ins01, g0ins01, l0ins01, n0ins01)	Subscription Server (EcSbSubServer) Event Server (EcSbEventServer)

Table 17.7-13 presents the steps required for recovering from a failure to retrieve data. The procedure is used to:

- make the initial check on the Science Data Server GUI.
- follow up with checks of the Science Data Server log file, Archive Server log file, and Request Manager log file.
- determine if a failure occurred during copying of the files to a staging area (and if so, whether there is sufficient staging space available).

If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 7 Launch the Science Data Server GUI (see Procedure 17.3.1 **Launching DSS GUIs**).
- 8 Click on the **System Requests** tab.
 - The **System Requests** window is displayed.
- 9 Examine the requests displayed in the **System Management Requests** field to determine if SDSRV received the acquire request.
 - If the number of request is large, the **Find** button and field below the **System Management Requests** field may be used to enter and search for information in the request, such as the Requester, or the **Filter . . .** button can be used to launch a **System Management Filter Requests** window to limit the number of entries that appear in the **System Management Requests** field.
- 10 On the SDSRV Server host (e.g., e0acs05, g0acs03, l0acs03, n0acs04), review the server logs **EcDsScienceDataServer.ALOG** and **EcDsScienceDataServerDebug.log** (refer to Procedure 17.7.2.1 **Checking Server Log Files**).
 - Examine the section of the log with entries near the time of the problem, looking for messages that indicate whether the relevant file was successfully acquired.
 - The **EcDsScienceDataServer.ALOG** file should contain entries identifying the file to be acquired by the ShortName of the corresponding ESDT; entries should be similar to the following:

```
PID : 29168:MsgLink :0 meaningfulname
:DsSrSessionExecuteRequestStart0
Msg: Request ID b5156038-03d3-11d3-8d16-c676e82eaa77:????:
executing:
DsSrRequest (1): DsShSciRequestImp: [ svr: ScienceDS, pri: NORMAL
domain: ]: (DsShSciCommandImp: service: INSERT num parameters: 3
category: Parameters are:
-UnnamedPL[SHORTNAME(AST_L1BT) VERSIONID(001)
--MAINGROUP[SHORTNAME(AST_L1BT) VERSIONID(001)
---
METAFILEGROUP[METADATAFILE(/home/cmops/data/SCAST_L1BT.0011279.met)]
---DATAFILEGROUP[DATAFILE(/home/cmops/data/tahoe-north-middle.MTA)]
---DATAFILEGROUP[DATAFILE(/home/cmops/data/tahoe-north-
middle.hdf)]]]
WC)
```

- The **EcDsScienceDataServerDebug.log** file should contain entries regarding the acquire activity. The following types of messages should be included in the log file:

```
05/06/99 12:52:01:
About to execute Statement: exec ProcInsertReqDomain 2205,
"UR:10:DsShESDT
UR:UR:15:DsShSciServerUR:13:[VTC:DSSDSRV]:20:SC:AST_L1BT.001:2201"
05/06/99 12:52:01:
About to execute Statement: ProcInsertAcquireCmd 2206, 2205, 3,
null, null, "tester", "FtpPush", "MAIL", "FILEFORMAT", null,
"jrattiga", "abc123", "t1dps04", "/home/jrattiga
/push", null, null
```

- If the ShortName does not appear in the file, with a timestamp corresponding to the time of the attempted acquire, SDSRV may not be running, or may not be communicating with other servers. Have the System Administrator or Operations Controller check to be sure the server is up and, if appropriate, resolve the problem (e.g., restart affected servers, execute **EcCsIdPingServers**, ensure that the **Name Server** is up in the mode being used and that its debug log reflects appropriate look-up activity by the application servers, mount points are intact, and database access is not impaired).
- If the log file does contain entries for the relevant ShortName, and indicates that two files (the file and its associated metadata file) are being distributed, SDSRV has completed its role in the acquire. Go to the next step.
- If the ALOG contains the ShortName, and also contains an error showing that the data file time stamp does not match the time stamp required by the acquire, the data file needs to be removed from the Science Data Server and reinserted.

– This is usually done using a script called DsDbCleanGranules.

- 11 To inspect the Archive Server log and Request Manager Server log for error messages associated with the acquire, on the Archive host (e.g., **e0drg11**, **g0drg01**, **l0drg01**, **n0drg01**), review the respective server logs (**EcDsStArchiveServerDebug.log**, **EcDsStRequestManagerServerDebug.log**); refer to Procedure 17.7.2.1 **Checking Server Log Files** and Procedure 17.7.2.2 **A Special Case: Checking the Request Manager Server Debug Log**.

- Examine the sections of the logs with entries near the time of the problem, looking for messages that indicate whether the Request Manager handled the request and whether the Archive Server log shows that the relevant file was successfully acquired.
- If the logs indicate that the relevant file was successfully acquired, go to the next step.
- If the file was not successfully acquired, it may be necessary to reboot AMASS (see Procedure 17.1.3 **Rebooting AMASS**) and investigate the possibility of read errors (see Procedure 17.7.4 **Diagnosing/Investigating Read Errors**).

- 12 To determine whether the file being acquired (or a link to it) and its associated metadata file arrived in the Data Distribution staging area, on the Distribution Server (e.g., **e0dis02**, **g0dis02**, **l0dis02**, **n0dis02**) type **cd /usr/ecs/<MODE>/CUSTOM/drp/<archivehost>/data/staging/cache** and then press the **Return/Enter** key.
- The working directory is changed to the specified directory.
- 13 Type **ls -lrt** and then press the **Return/Enter** key.
- The contents of the directory are displayed.
- 14 Review the listing to determine whether the relevant file and its metadata file arrived in the staging area.
- The display should contain entries similar to the following:

```
lrwxrwxr-x  1 cmshared cmshared  75 Apr 26 12:52
L7CPF19980518_19980518.01 ->
/usr/ecs/TS1/CUSTOM/drp/raven/data/staging/cache/:SC:L7CPF.001:1427:
1.ASCII
-rw-rw-rw-  1 cmshared cmshared  14802 Apr 26 12:52
SCL7CPF.0011427.met
-rw-rw-r--  1 cmshared cmshared    111 Apr 26 13:01
staging.disk.filename.list
-rw-rw-r--  1 cmshared cmshared  2044 Apr 26 13:01
PACKING.LST.115124935248431
```
 - If the relevant files were not successfully staged, the staging log files may reveal the cause; go to Step 9.
 - If the relevant files were successfully staged, an acquire failure could be a result of problems with related servers or software (see Table 1). Have the System Administrator or Operations Controller ensure that the necessary hosts and servers are up.
- 15 To inspect the Staging Disk log for error messages associated with the acquire, on the APC Server host (e.g., **e0acg11**, **g0acg01**, **l0acg02**, **n0acg01**), review the server logs (e.g., **EcDsStStagingDiskServerDebug.log**; **EcDsStCacheManagerServerDebug.log**); refer to Procedure 17.7.2.1 **Checking Server Log Files**.
- Examine the section of each log with entries near the time of the problem, looking for messages that indicate whether the relevant files were successfully staged.
 - If the relevant files were not successfully staged, the cause may be a lack of space in the staging area; go to Step 10.
 - If the relevant files were successfully staged, an acquire failure could be a result of problems with related servers or software (see Table 1). Have the System Administrator or Operations Controller ensure that the necessary hosts and servers are up.

16 To check the space available in the staging area, on the Distribution Server (e.g., **e0dis02**, **g0dis02**, **l0dis02**, **n0dis02**) type **cd /usr/ecs/<mode>/CUSTOM/drp/<archivehost>/data** and then press the **Return/Enter** key.

- The working directory is changed to the specified directory.

17 Type **df -k .** (be sure to include the ".") and then press the **Return/Enter** key.

- The filesystem, staging disk space capacity in kbytes, amount used, amount available, and percent of capacity are displayed, as in the following example:

```
Filesystem          kbytes    used   avail capacity  Mounted on
t1drg01:/usr/ecs/TS1/CUSTOM/drp/t1drg01/data
                    225209856 173253056 51956800    77%
/data1/ecs/TS1/CUSTOM/drp/t1drg01/data
```

- If there is not adequate space for staging the relevant files, it will be necessary to free up additional space (e.g., by purging expired files from cache).

Table 17.7-13. Handling a Data Acquire Failure

Step	What to Do	Action to Take
1	Launch the Science Data Server GUI	use Procedure 17.3.1
2	Select the System Requests tab	single-click
3	Review list of System Management Requests to determine if SDSRV receive the acquire request	read text
4	Review SDSRV logs for evidence of acquire or errors: a. EcDsScienceDataServer.ALOG b. EcDsScienceDataServerDebug.log	a. use Procedure 17.7.2.1 b. use Procedure 17.7.2.1
5	Review server logs for acquire error messages: a. EcDsStArchiveServerDebug.log b. EcDsStRequestManagerServerDebug.log	a. use Procedure 17.7.2.1 b. use Procedure 17.7.2.2
6	(On the Distribution Server host) cd /usr/ecs/<MODE>/CUSTOM/drp/<archivehost>/data/staging/cache	enter text; press Return/Enter
7	ls -lrt	enter text; press Return/Enter
8	Review listing for relevant file and metadata	read text
9	(On APC Server host) Review staging disk logs for acquire error messages: a. EcDsStStagingDiskServerDebug.log b. EcDsStCacheManagerServerDebug.log	a. use Procedure 17.7.2.1 b. use Procedure 17.7.2.1
10	(On the Distribution Server host) cd /usr/ecs/<mode>/CUSTOM/drp/<archivehost>/data	enter text; press Return/Enter
11	dr -k .	enter text; press Return/Enter

17.7.3 Diagnosing/Investigating Write Errors

Although write errors to the archive should be infrequent, there are some circumstances under which they may occur. Associated error messages should appear in a relevant log file (e.g., on a Sun, in **/var/adm/messages**; on an SGI, in **/var/adm/SYSLOG**). Causes of write errors may include the following:

- *AMASS off line* -- software captures and logs the error because the directory that is being written to does not exist. However, the nature of the write error is not detected.
- *Directory does not exist* -- if there is an attempt to write to a directory that does not exist, even if AMASS is on line, the result is a write error.
- *All drives off line* -- write requests are accepted until cache space fills up, which stops further data transfer.
- *Volume off line or no media associated with the directory* - causes a write error that is detectable by the software. An I/O error is recorded in the relevant log file.

- *AMASS: media write failure* -- causes the drive to go off-line and the media volume to go off-line as well. The error is written to the relevant log file. No error is detected by the application software. The operator can execute the command **/usr/amass/bin/drivelist** to see which drive has been put off-line.

Table 17.7- 14 presents the general steps required for diagnosing/investigating write errors. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01) as **amass**.
- 2 To verify that AMASS is running and active, type **/usr/amass/bin/amassstat -c** and then press the **Return/Enter** key.
 - The message **FILESYSTEM IS ACTIVE** is displayed.
 - If the message **FILESYSTEM IS INACTIVE** is displayed, AMASS is inactive (off line); to put AMASS back on line, type **/usr/amass/bin/amassstat -a** and press the **Return/Enter** key, entering **y** in answer to the displayed confirmation question **Do you want to change the status? [y - n]**. If this does not activate AMASS, it may be necessary to reboot it (refer to Procedure 17.1.3 **Rebooting AMASS**).
- 3 Display AMASS I/O activity (refer to Procedure 17.7.1.2 **Using sysperf to Display the Status of AMASS I/O Activity**).
 - If the returned output indicates no free cache blocks, this may indicate that all drives are off line and that as a result all available cache space is filled, stopping write actions. In that case, it may be necessary to use **ctrl-c** to cancel pending I/O requests.
 - If there are no free cache blocks, go to Step 4; otherwise, go to Step 5.
- 4 Type **/usr/amass/bin/drivelist** and press the **Return/Enter** key.
 - AMASS returns the status of drives in format similar to the following:

DRIVE	JUKE	STATUS	ERRORS
1	1	A	0
2	1	A	0
3	1	I	0
4	1	A	0
1	2	I	0
2	2	A	0
1	3	A	0
2	3	A	0
3	3	A	0

/usr/amass/bin/drivelist: 9 drives configured in this system

- 5 Use the **amass_log** script to display and examine AMASS errors (refer to Procedure 17.7.1.4 **Using the *amass_log* Script to Display AMASS Errors**).
 - If the returned error message(s) identify critical problems that prevent AMASS from functioning correctly, follow the corrective guidance specified in the *System Administrator's Guide* for the specific error(s).
 - If there are no hardware problems and there have not been repeated attempts to activate the drive(s) that are off line, reactivate the drive(s) (see Step 6).
- 6 For each off-line drive, type **/usr/amass/bin/drivestat -a drivenumber [jukeboxnumber]** and press the **Return/Enter** key.
 - AMASS places on line the drive specified by **drivenumber** in the jukebox specified by **jukeboxnumber**.
- 7 Check the **el_ETAC** log for evidence of problems in communication between AMASS and ACSLS (refer to procedure 17.7.2.3 **Checking the *el_ETAC* Log**).
 - If there is evidence of communication problems, investigate and correct any lack of synchrony between AMASS and ACSLS (refer to Procedure 17.7.1.6 **Using *mediamove* to Establish Synchrony Between *quedisplay* and *medialist***).

Table 17.7- 14. Diagnosing/Investigating Write Errors

Step	What to Do	Action to Take
1	Log in to the FSMS host as amass	enter text; press Return/Enter
2	/usr/amass/bin/amassstat -c	enter text; press Return/Enter
3	Display status of AMASS I/O activity	use Procedure 17.7.1.2
4	(If no free cache blocks) /usr/amass/bin/drivelist	enter text; press Return/Enter
5	Run amass_log script	use Procedure 17.7.1.4
6	/usr/amass/bin/drivestat -a drivenumber [jukeboxnumber]	enter text; press Return/Enter
7	Check el_ETAC log	use Procedure 17.7.2.3

17.7.4 Diagnosing/Investigating Read Errors

When a read error is encountered by AMASS, both the drive and the volume (tape) are taken off line. The application is notified of the read failure. The Archive Server logs an error message when the read failure is returned. The log message includes the name of the file, the secondary path for the file, the checksum for the file, and a reason for the failure. If the reason for failure is a checksum mismatch on retrieval, then the file must be restored. If the reason for failure indicates the media was off line, then further investigation is warranted to determine why the tape was off line. Off-line status can be the result of a write error, a read error on the file, or a read error on another file that caused AMASS to take the tape off line, thus making other reads fail. For a tape that is off line, or for a tape and drive that are off line together, one possibility is that the tape is damaged. Damage may be confirmed by visual inspection or, more likely, the

need to have vendor maintenance remove the media from the drive. Any requests for files on that tape fail or are served from backup. It is important that the list of files that is created for restoring a tape from backup be kept and searched when new files are reported missing. This should reduce the number of times that certain recovery procedures have to be performed (see Procedure 17.6.2.3 **Manual Data Recovery from Damaged Cartridge**).

Table 17.7-15 presents the steps required for diagnosing/investigating read errors observed during operations (e.g., appearing in the Archive Server debug log, or appearing as an I/O error message at the command line during an attempt to copy a file from an archive volume). If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01) as **amass**.
- 2 Examine the Archive Server debug log **EcDsStArchiveServerDebug.log** (refer to Procedure 17.7.2.1 **Checking Server Log Files**) for error messages associated with the read failure.
 - Examine the sections of the log with entries near the time of the failure, looking for messages that indicate read failure. It may be useful to search the log for occurrences of the word **fail** (while viewing the log with **pg**, **view**, **vi**, or other viewing/editing tool, type **/fail** and press the **Return/Enter** key).
 - From the failure information in the log, note the name of the file, the secondary path for the file, the checksum for the file, and the reason for the failure.
 - If the reason for the failure is specified as a checksum mismatch on retrieval, go to Step 3. If the reason indicates media being off line, go to Step 4.
- 3 Restore the corrupted file (refer to Procedures 17.6.2.1 **Manual Data Recovery from Local Backup Tapes**, 17.6.2.2 **Manual Data Recovery from Offsite Backup Tapes**, and 17.6.2.3 **Manual Data Recovery from Damaged Cartridge** as appropriate).
- 4 Type **dirfilelist path**, where **path** is the full path name of the directory location of the file on which the read error occurred (e.g., **/dss_stk1/OPS/modl0**), and then press the **Return/Enter** key.
 - AMASS returns a listing of the files in the directory, listing for each one the volume on which it is stored and its logical block address. Note the volume number for the file on which the read error occurred.
- 5 Use the **vollist** command to display data for the volume identified in Step 4 (refer to Procedure 17.7.1.3 **Using vollist to Display Volume Data**).
 - AMASS displays data for the specified volume; if the volume is off line (has **O** displayed in the **FLAGS** column of the output), place it on line using the command **volstat -a** and pressing the **Return/Enter** key.

- 6 Use the **amass_log** script to display and examine AMASS errors (refer to Procedure 17.7.1.4 **Using the *amass_log* Script to Display AMASS Errors**).
 - If the returned error message(s) identify critical problems that prevent AMASS from functioning correctly, follow the corrective guidance specified in the *System Administrator's Guide* for the specific error(s).
 - If there are no hardware problems and there have not been repeated attempts to activate the drive(s) that are off line, reactivate the drive(s) (see Step 7).
- 7 For each off-line drive, type **/usr/amass/bin/drivestat -a drivenumber [jukeboxnumber]** and press the **Return/Enter** key.
 - AMASS places on line the drive specified by **drivenumber** in the jukebox specified by **jukeboxnumber**.
- 8 Check the **el_ETAC** log for evidence of problems in communication between AMASS and ACSLS (refer to procedure 17.7.2.3 **Checking the *el_ETAC* Log**).
 - If there is evidence of communication problems, investigate and correct any lack of synchrony between AMASS and ACSLS (refer to Procedure 17.7.1.6 **Using *mediamove* to Establish Synchrony Between *quedisplay* and *medialist***).

Table 17.7-15. Diagnosing/Investigating Read Errors

Step	What to Do	Action to Take
1	Log in to the FSMS host as amass	enter text; press Return/Enter
2	Examine the Archive Server debug log	use Procedure 17.7.2.1
3	Restore any corrupted file	as appropriate, use Procedure 17.6.2.1, 17.6.2.2, or 17.6.2.3
4	dirfilelist path	enter text; press Return/Enter
5	vollist nnn	use Procedure 17.7.1.3
6	Run amass_log script	use Procedure 17.7.1.4
7	/usr/amass/bin/drivestat -a drivenumber [jukeboxnumber]	enter text; press Return/Enter
8	Check el_ETAC log	use Procedure 17.7.2.3

17.8 ACSLS Procedures

For the StorageTek Powderhorn, direct control of the tape storage and handling operations is managed by the *Automated Cartridge System Library Software (ACSL)*. Full guidance for using ACSLS is provided in the *Automated Cartridge System Library Software System Administrator's Guide*. Table 17.8-1 lists the commands covered in that *Guide*.

Table 17.8-1. ACSLS Command Reference

Command	Function
audit	Creates or updates the database inventory of the volumes in a library component.
cancel	Cancels a current or pending request.
clear lock	Removes all active and pending locks on transports or volumes
define pool	Creates or modifies scratch pools.
delete pool	Deletes empty scratch pools.
dismount	Dismounts a volume.
eject	Ejects one or more volumes from the Automated Cartridge System (ACS).
enter	Sets a Cartridge Access Port (CAP) to enter mode.
idle	Stops ACSLS from processing new requests.
lock	Locks (dedicates) a volume or transport to a user.
logoff	Exits the command processor.
mount	Mounts a data or scratch volume.
query	Displays the status of a library component.
set	Sets various attributes of different library components.
show	Displays your lock ID or user ID.
start	Starts ACSLS request processing.
unlock	Removes active locks on volumes or transports.
vary	Changes the state of an ACS, LSM, CAP, transport, or port.
venter	Enters one or more volumes with missing or unreadable labels into the ACS.

ACSLS commands use the following general syntax:

command type_identifier state [options]

where **type_identifier** is the ACS component and its identifier (these are listed in the *System Administrator's Guide*), **state** is a device state for the **vary** command only, and **options** are command options (these are specified for each command in the *System Administrator's Guide*). The two most useful commands in ACSLS are **query** and **vary**. Other frequently used commands are **enter** and **eject**, for inserting and removing cartridges, respectively. ACSLS does not have an online help facility, but if you enter a command (e.g., **vary**), it will prompt you for the parameters.

There are also several utilities provided with ACSLS. These are listed with their functions in Table 17.8-2.

Table 17.8-2. ACSLS Utilities

Utility	Function
bdb.acsss	Backs up the ACSLS database.
kill.acsss	Terminates ACSLS.
rc.acsss	Starts and recovers ACSLS.
rdb.acsss	Restores the ACSLS database.
Volrpt	Creates a volume report.
db_command	Starts or stops the Oracle database.

To control and interact with ACSLS, you use the following user IDs:

- **acssa** lets you enter ACSLS commands from a command processor window.
- **acsss** lets you run ACSLS utilities from the UNIX command line prompt.

It is typical to log in as both user IDs to permit entering both ACSLS utilities and commands. You can, however, open a command processor window from the **acsss** user ID if you prefer to work from a single user ID. The *System Administrator's Guide* provides full details. Table 17.8-3 provides an Activity Checklist for major ACSLS procedures addressed in this section.

Table 17.8-3. Activity Checklist for ACSLS Procedures

Order	Role	Task	Section	Complete?
1	Archive Manager	Entering the Archive after AMASS is Started	(P) 17.8.1	
2	Archive Manager	Backing up the ACSLS Database	(P) 17.8.2	
3	Archive Manager	Restoring the ACSLS Database	(P) 17.8.3	
4	Archive Manager	Checking Cleaning Cartridges	(P) 17.8.4	

17.8.1 Entering the Archive After AMASS is Started

There are circumstances in which it may be necessary to enter the archive after AMASS is started. For example, there may be a requirement for maintenance that necessitates access to the robot or other area inside the Powderhorn. Another example is that it may sometime be desirable to bypass the Cartridge Access Port (CAP) when inserting tape cartridges, if there is a need to perform bulk loading of a large number of tapes, although usually this would be limited to initial loading of the volumes. Table 17.8-4 presents the steps required for entering the archive after AMASS has started. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 At the host for ACSLS (e.g., **e0drs03**, **g0drs03**, **l0drs02**, **n0drs03**), log in using the **acssa** user ID and password.
 - The **acssa** command-process window is displayed with the **ACSSA>** prompt.
- 2 Type **vary lsm 0,0 offline** and then press the **Return/Enter** key.
 - The access port is unlocked (audible unlatching sound).
- 3 Use the key to unlatch and open the access door.
 - A red **DO NOT ENTER** warning is visible inside the enclosure.

Warning

If it is necessary to enter the STK Powderhorn after AMASS is started, it is necessary to perform the following step to avoid hazard and ensure safety of personnel and equipment.

- 4 Remove the key from the door to ensure that no one inadvertently locks the enclosure with someone inside.
 - The red **DO NOT ENTER** warning is extinguished and a green **ENTER** message is displayed inside the enclosure.
- 5 Upon leaving the enclosed area, insert the key in the access door and latch the door.
 - The LED display indicates that the door is locked.
- 6 At the ACSLS host, type **vary lsm 0,0 online** and then press the **Return/Enter** key. After a few seconds, the archive robots execute an initialization sequence and the LSM is back online.

Table 17.8-4. Entering the Archive after AMASS is Started

Step	What to Do	Action to Take
1	Log in to the ACSLS host as acssa	enter text; press Return/Enter
2	vary lsm 0,0 offline	enter text; press Return/Enter
3	Unlock and open access door	operate lock to unlatch with key
4	Remove key from unlatched door	remove key
5	Insert key and latch door	operate lock to latch with key
6	vary lsm 0,0 online	enter text; press Return/Enter

17.8.2 Backing Up the ACSLS Database

ACSLs provides the **bdb.acsss** utility to back up the database. It is advisable to run this utility when there has been a change in the archive volume structure (e.g., upon addition or removal of volumes). In the event of database loss, it is possible to re-create the database even if there is no backup available, by using the **audit** command to inventory the archive. However, for a large storage facility, creating the database this way may take several hours. If there is a backup available, the database can be restored easily and quickly (refer to Procedure 17.8.3). Table 17.8-5 presents the steps required for backing up the ACSLS database. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 At the host for ACSLS (e.g., **e0drs03**, **g0drs03**, **l0drs02**, **n0drs03**), log in using the **acsss** user ID and password.
 - The **acsss** command-process window is displayed with the **ACSSS>** prompt.
- 2 Ensure that there is a tape in the backup drive (device **dev/rmt/0**), a streaming tape drive attached to each ACSLS workstation.
- 3 Type **bdb.acsss**, and then press the **Return/Enter** key.
 - If you enter **bdb.acsss** with no options, the backup utility defaults to the default tape device attached and configured to the ACSLS server.
 - The system displays the following message.


```
Check tape device (/dev/rmt/0) to make sure you have a tape in the
tape drive.

[ Hit RETURN to continue or Ctrl-C to exit ]
```
- 4 Press the **Return/Enter** key.
 - The **bdb.acsss** utility backs up the ACSLS database and miscellaneous library resource files.

Table 17.8-5. Backing Up the ACSLS Database

Step	What to Do	Action to Take
1	Log in to the ACSLS host as acsss	enter text; press Return/Enter
2	Ensure there is a tape in the backup drive	Mount tape
3	bdb.acsss	enter text; press Return/Enter
4	Return/Enter	press Return/Enter

17.8.3 Restoring the ACSLS Database

ACSLs provides the **rdb.acsss** utility to restore the database in case of severe disk or data problems. If you have made regular backups, it should be possible to restore the database with little or no loss of data. Restoring the database is likely to be necessary if there has been a system crash, or if the database can not be started or has a physical or logical error. Table 17.8-6 presents the steps required for restoring the ACSLS database. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 At the host for ACSLS (e.g., **e0drs03**, **g0drs03**, **l0drs02**, **n0drs03**), log in using the **acsss** user ID and password.
 - The **acsss** command-process window is displayed with the **ACSSS>** prompt.

- 2 Load the restore tape into the backup drive.
- 3 Type **rdb.acsss**, and then press the **Return/Enter** key.
 - If you enter **bdb.acsss** with no options, the backup utility defaults to the default tape device attached and configured to the ACSLS server.
 - The system displays the following message.


```
Check tape device (/dev/rmt/0) to make sure you have a tape in the
tape drive.

[ Hit RETURN to continue or Ctrl-C to exit ]
```
- 4 Press the **Return/Enter** key.
 - The rdb.acsss utility restores the ACSLS database and miscellaneous library resource files.

Table 17.8-6. Restoring the ACSLS Database

Step	What to Do	Action to Take
1	Log in to the ACSLS host as acsss	enter text; press Return/Enter
2	Load the restore tape into the backup drive	Mount tape
3	rdb.acsss	enter text; press Return/Enter
4	Return/Enter	press Return/Enter

17.8.4 Checking Cleaning Cartridges

The Automated Cartridge System Library Software (ACSL) schedules and implements routine cleaning of the system tape drives after a set usage time tracked by the software, using cleaning volumes from a cleaning volume group designated for that purpose. The ACSLS software also tracks the number of times a cleaning tape is used, and will not use a cleaning tape that has been used the maximum set number of times (usually set at 100 for the 9940 drives). It is the responsibility of the Archive Manager to monitor cleaning tape usage periodically, to ensure that usable cleaning tapes remain available to the system.

Table 17.8-7 presents the steps required to check cleaning cartridges for usage status. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 At the host for ACSLS (e.g., **e0drs03**, **g0drs03**, **l0drs02**, **n0drs03**), log in using the **acssa** user ID and password.
 - The **acssa** command-process window is displayed with the **ACSSA>** prompt.
- 2 Type **query clean all**, and press the **Return/Enter** key.
 - **Note:** The command may be abbreviated to **qu cl a**.

- ACSLS displays information on the status of the cleaning volumes in format similar to the following:

```

2001-10-04 08:50:54      Cleaning Cartridge Status
Identifier Home Location  Max Usage  Current Usage  Status  Type
9840C1      0, 0, 3, 2, 2  100        38             home    STK1U
9840C2      0, 0,13, 1, 3  100         0             home    STK1U
9940C1      0, 0, 1, 4,19  100         7             home    STK2W

```

- **Note:** If it is desirable or necessary to change the maximum number of uses permitted for a cleaning volume, the change can be accomplished with the command **set clean <max_usage> <vol_id>** where *max_usage* (e.g. 100) is the maximum number of uses for that volume and *vol_id* is the volume id of that cleaning cartridge.

Table 17.8-7. Checking Cleaning Cartridges

Step	What to Do	Action to Take
1	Log in as acssa	enter text; press Return/Enter
2	query clean all (or qu cl a)	press Return/Enter

17.9 Using the AMASS Graphical User Interface (GUI)

AMASS offers a Graphical User Interface (GUI) called the AMASS Administration Window (AAWin) through which operators can administer volumes and volume groups that are managed by AMASS. AAWin provides a point-and-click interface for identifying volumes their groups, and their configurable parameters. The AAWin main window is composed of a menu bar, a large middle section called the *workroom*, a utility bar at the right with icons for a trash can, a volume group, and a volume, and a status bar at the bottom with indicator “lights” that represent the current status of AMASS. Selecting the volume icon on the utility bar populates the workroom with icons for volumes. Moving the cursor over one of the icons results in the appearance of volume-related information in a pop-up display.

For large storage facilities, command-line interactions are likely to be faster and more responsive than interactions with the AMASS GUI. Therefore it is unlikely that extensive use of the GUI will be applied for most archive operations. However, it may be useful to have the GUI open for monitoring and easy access to volume information (refer to Procedure 17.9.1 **Launching the AMASS GUI and Viewing Volume Group and Volume Information**). Table 17.9-1 provides an activity checklist for some procedures that may be accomplished using the AMASS GUI.

Table 17.9-1. Activity Checklist for Using the AMASS GUI

Order	Role	Task	Section	Complete?
1	Archive Manager	Launching the AMASS GUI and Viewing Volume Group and Volume Information	(P) 17.9.1	
2	Archive Manager	Using the AMASS GUI to Modify a Volume Group	(P) 17.9.2	
3	Archive Manager	Using the AMASS GUI to Modify a Volume	(P) 17.9.3	

17.9.1 Launching the AMASS GUI and Viewing Volume Group and Volume Information

Table 17.9-2 presents the steps required to launch the AMASS GUI and view information about volume groups and volumes in the archive. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in as **amass** at the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01).
- 2 Type **/usr/amass/bin/aawin** and then press the **Return/Enter** key.
 - The AMASS GUI main window is displayed.
- 3 Click on the **View by Volume Groups** button (middle button at the right of the *workroom*).
 - The *workroom* is populated by icons for volume groups.
 - The **Block List** window is displayed; it is a vertically scrolled list of blocks of items (in this case, volume groups). The *workroom* can display up to 256 icons; the **Block List** window provides access to additional items in blocks of 256.
- 4 Move the cursor over one of the icons for a volume group.
 - A pop-up display shows data for the volume group (**Volume Group, Volumes in Group, Free Space, Dead Space, Error Count**).
- 5 Click on the **View by Volumes** button (at the bottom right side of the *workroom*).
 - The *workroom* is populated by icons for volumes.
 - The **Block List** window is also displayed; it is a vertically scrolled list of blocks of items (in this case, volumes).
- 6 Move the cursor over one of the icons for a volume.
 - A pop-up display shows data for the volume group (**Volume, Library, Slot, Volume Group, Volume Status, Volume Label**).

Table 17.9-2. Launching the AMASS GUI and Viewing Volume Group and Volume Information

Step	What to Do	Action to Take
1	Log in as amass at FSMS host	enter text; press Return/Enter
2	aawin	enter text; press Return/Enter
3	Select View by Volume Groups button	single-click
4	Move cursor over icon for a volume group	hold left mouse button and drag
5	Select View by Volumes button	single-click
6	Move cursor over icon for a volume	hold left mouse button and drag

17.9.2 Using the AMASS GUI to Modify a Volume Group

The **Modify a VG** window is opened by selecting **Modify a Volume Group** from the **Tasks** menu. The window is used to modify the characteristics of a volume group. The top portion of the window (not modifiable) lists root directories already configured for a volume group. The

middle portion of the window permits adding directories to the list of root directories for the specified volume group. The third major portion of the window, near the bottom, contains indicators of the status of the volume group and buttons for selecting a volume group, as well as buttons across the very bottom of the window for accepting or canceling the modifications. (*Note:* The **Modify a VG** window also is opened if you have the *workroom* populated with volume group icons and you click on one of them. However, in this case you may only modify the volume group on which you clicked; the bottom of the window will not display buttons for selecting a volume group.)

As an example of using the GUI to modify a volume group, it is possible to assign a new root directory in the AMASS file system to a volume group. This requires first creating the directory and then using the GUI to assign it to a volume group. Table 17.9-3 presents the steps required for using the AMASS GUI to modify a volume group, in this example to assign a new root directory in the AMASS file system to a volume group. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the AMASS GUI (refer to Procedure 17.9.1 **Launching the AMASS GUI and Viewing Volume Group and Volume Information**).
- 2 Open a second terminal window (other than the one used to launch the AMASS GUI).
- 3 In the second terminal window, log in as **amass** at the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01).
- 4 To change to the **dss_amass** directory, type **cd /dss_amass**, and then press the **Return/Enter** key.
- 5 To create an empty directory with path **/dss_amass/newdir/**, where *newdir* is the name of the new directory to be created and assigned to the volume group, type **mkdir newdir**, and then press the **Return/Enter** key.
- 6 On the AMASS GUI main window, click on the **View by Volume Groups** button (middle button at the right of the *workroom*).
 - The *workroom* is populated by icons for volume groups.
 - The **Block List** window is also displayed; it is a vertically scrolled list of blocks of items (in this case, volume groups).
- 7 Follow menu path **Tasks→Modify a Volume Group**.
 - The **Modify a VG** window is displayed, showing data for Volume Group 0001.
- 8 In the area for choosing a volume group, near the bottom of the window, use the buttons to set the number displayed in the **Volume Group** field to the desired volume group.
 - A click on the right-pointing arrow button or the left-pointing arrow button respectively increases or decreases the number by one. Buttons below the arrow

buttons may be used to increase or decrease the number in multiples of 100 or 1000, as indicated on the buttons.

- 9 When the **Volume Group** field displays the number of the desired volume group, click on the **Fetch** button.
 - The list of root directories already configured for the selected volume group is displayed in the **Existing Root Directories** field.
 - The status indicators show the status of the selected volume group.
- 10 Click on the **File/Directory Selection** button (leftmost button after the label **Root Directories to Add**, with folder icon).
 - A **File Selection** filter window is displayed.
- 11 In the **File Selection** filter window, click on the **Filter** button.
 - The **Filter** field displays **/usr/amass/***, and directories and files are displayed in the **Directories** and **Files** windows, respectively.
- 12 Use the **Filter** button and selection of directories in the **Directories** window to display **/dss_amass/newdir/** in the **Selection** field.
 - The **Selection** field displays **/dss_amass/newdir/**.
- 13 In the **File Selection** filter window, click the **OK** button.
 - The **Root Directories to Add** field of the **Modify a VG** window displays **/dss_amass/newdir/**.
- 14 To examine the capability to edit the list of Directories to Add, click on the entry **/dss_amass/newdir/** to highlight it in the **Root Directories to Add**, then click on the **Remove a File/Directory from List** button (middle button after the label **Root Directories to Add**, with folder icon crossed out with a red line).
 - The entry **/dss_amass/newdir/** is removed from the **Root Directories to Add** field.
- 15 Repeat steps 11 - 13 to restore the entry **/dss_amass/newdir/** to the **Root Directories to Add** field.
 - The **Root Directories to Add** field of the **Modify a VG** window displays **/dss_amass/newdir/**.
- 16 In the **Modify a VG** window, click on the **Accept** button at the bottom of the window.
 - The entry **/dss_amass/newdir/** is removed from the **Root Directories to Add** field and appears in the **Existing Root Directories** field.
 - The **Modify a VG** window is closed.

Table 17.9-3. Using the AMASS GUI to Modify a Volume Group

Step	What to Do	Action to Take
1	Launch the AMASS GUI	use Procedure 17.9.1
2	Open a second terminal window	UNIX command
3	Log in as amass at FSMS host	enter text; press Return/Enter
4	cd /dss_amass	enter text; press Return/Enter
5	mkdir newdir	enter text; press Return/Enter
6	Select View by Volume Groups button	single-click
7	Follow menu path Tasks→Modify a Volume Group	menu selection
8	Set Volume Group	click on right or left arrow
9	Select Fetch button	single-click
10	Select File/Directory Selection button	single-click
11	Select Filter button	single-click
12	Display /dss_amass/newdir/ in Selection field	use Filter button and seelction
13	In File Selection filter window, select OK button	single-click
14	Remove /dss_amass/newdir/ from list	highlight and click
15	Restore /dss_amass/newdir/ to list	highlight and click
16	Select Accept button	single-click

17.9.3 Modify a Volume

The **Modify a Volume** window is opened by selecting **Modify a Volume** from the **Tasks** menu. The window is used to modify the characteristics of a volume. The right side of the window shows the current set of statistics and configuration information (not modifiable) for the volume listed in the **Volume** field on the left side of the window (the **Volume** field looks like a button, but if you click on it, a “spinbox” is displayed, with arrow buttons permitting increases or decreases to the volume number, and buttons at the bottom to **Accept** or **Cancel** the change; accepting the change closes the spinbox, displays the new number in the **Volume** field, and displays data for that volume). The left side of the **Modify a Volume** window provides access to modifiable characteristics of the volume. Changes made to the buttons and fields in the window do not take effect until the **Accept** button at the bottom of the window is clicked. (**Note:** The **Modify a Volume** window also is opened if you have the *workroom* populated with volume icons and you click on one of them. However, in this case you may only modify the volume on which you clicked; the **Volume** field does not look like a button and may not be changed.)

There are six fields that can be edited for a given volume:

1. The first is a button for setting the **Volume Group**. Clicking the button opens a spinbox for selecting the volume group to which the volume is to be assigned.
2. Below the Volume Group button is an **Online/Offline** indicator light with label. Clicking on the indicator toggles its state and updates the text field (label) next to it.
3. Below the Online/Offline indicator is an **Active/Inactive** indicator light with label. Clicking on the indicator toggles its state and updates the text field (label) next to it.
4. Next is a **Format Request** option button permitting selection of a formatting option for the volume.

5. Next is the **Block Size** field, applicable only to tape libraries when a format is requested to be done on the volume. This field requires a numeric value, which should be a multiple of 16384.
6. The last modifiable field is a text field for specifying the volume label.

Table 17.9-4 presents the steps required for using the AMASS GUI to modify a volume. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the AMASS GUI (refer to Procedure 17.9.1 **Launching the AMASS GUI and Viewing Volume Group and Volume Information**).
- 2 On the AMASS GUI main window, click on the **View by Volume Groups** button (middle button at the right of the *workroom*).
 - The *workroom* is populated by icons for volume groups.
 - The **Block List** window is also displayed; it is a vertically scrolled list of blocks of items (in this case, volume groups).
- 3 Click on the icon for a desired volume group.
 - The *workroom* is populated with icons for the volumes in the selected volume group, and the **Modify a VG** window is displayed, showing data for the selected volume group.
- 4 Click on the icon for the volume to be modified.
 - The **Modify a VG** window is closed and the **Modify a Volume** window is displayed, showing data for the selected volume.
- 5 If it is desired to change the volume group to which the volume is assigned, note the **Volume Group** number indicated on the **Volume Group** button, and then click on the button.
 - A spinbox is displayed showing the **Volume Group** number, with right-pointing and left-pointing arrow buttons respectively to increase or decrease the number.
- 6 Use the arrow buttons to change the **Volume Group** number, and then click on the **Accept** button in the spinbox.
 - The spinbox is closed and the new number appears in the **Modify a Volume** window as the **Volume Group** number.
- 7 To change the status of a volume indicated to be **Online**, click on the **Active/Inactive** indicator.
 - The color and label of the **Active/Inactive** indicator toggle.

- 8 To change the status of a volume indicated to be **Inactive**, click on the **Online/Offline** indicator.
 - The color and label of the **Online/Offline** indicator toggle.
- 9 Click on the **Format Request** option button.
 - A pop-up option menu is displayed for selection of **Yes** or **No**, and when one of those options is clicked, the indicated choice is displayed on the option button.
- 10 Use the mouse to move the cursor to the **Block Size** field.
 - A blinking cursor appears in the **Block Size** field.
- 11 Use the keyboard to enter or change the value in the **Block Size** field.
 - The entered data appear in the **Block Size** field.
- 12 Use the mouse to move the cursor to the **Volume Label** field.
 - A blinking cursor appears in the **Volume Label** field.
- 13 Use the keyboard to enter or change the value in the **Volume Label** field.
 - The entered data appear in the **Volume Label** field.
- 14 If you wish to cancel any request for changes to the volume, click on the **Cancel** button at the bottom of the window. If you wish to accept the changes, click on the **Accept** button at the bottom of the window.
 - When you click the **Accept** button, *AAWin* attempts to make the requested changes. For most changes, specifically changes to **Online/Offline** and **Active/Inactive** status, the requested **Volume Group** for the volume, and the **Volume Label**, the changes are made immediately. But if a format has been requested, then the **Online/Offline** and **Active/Inactive** status changes are not applied immediately. Instead, the requests for these status changes and the format changes are passed to the *AAWin Scheduler* daemon for processing. Changes made by the **Scheduler** occur when the job is processed, which depends on how many other jobs are currently scheduled.

Table 17.9-4. Using the AMASS GUI to Modify a Volume

Step	What to Do	Action to Take
1	Launch the AMASS GUI	use Procedure 17.9.1
2	Select View by Volume Groups button	single-click
3	Select desired volume group	single-click
4	Select volume to be modified	single-click
5	To change volume group assignment, select Volume Group button	single-click
6	Set Volume Group number	click arrows on "spinbox"
7	Toggle Active/Inactive status of online volume	single-click

8	Toggle Online/Offline status of inactive volume	single-click
9	Select Format Request button	single-click
10	Move cursor to Block Size field	single-click
11	Enter Block Size	type text
12	Move cursor to Volume Label field	single-click
13	Enter Volume Label	type text
14	Select Cancel button or Accept button	single-click

17.10 Data Pool Maintenance Tasks

Archive and/or engineering support personnel are directly involved in Data Pool monitoring and maintenance, and support User Services and/or Science Data Specialists in managing the content and retention of data in the Data Pool. A major tool used for these functions is the Data Pool Maintenance (DPM) GUI. There are also scripts and utilities for specific maintenance and monitoring functions, and the Spatial Subscription Server (NBSRV) GUI can be used for some maintenance and monitoring functions. Table 17.10-1 provides an activity checklist for Data Pool procedures that are accomplished using the DPM GUI and Data Pool utility scripts. See Chapter 19 for procedures using the NBSRV GUI (Launching the NBSRV GUI; Use the NBSRV GUI to List and View Subscriptions in the NBSRV Database; Use the NBSRV GUI to Extend the Period of Retention in a Data Pool Insert Subscription; Use the NBSRV GUI to View the Acquire and Notification Actions Being Processed; Use the NBSRV GUI to View Statistics on NBSRV Processing of Events and Actions).

Table 17.10-1. Activity Checklist for Data Pool Maintenance Tasks

Order	Role	Task	Section	Complete?
1	Archive Manager/ Support Engineer	Launch the Data Pool Maintenance (DPM) GUI	(P) 17.10.1	
2	Archive Manager/ Support Engineer	Using the DPM GUI to Monitor Data Pool Active Insert Processes	(P) 17.10.2	
3	Archive Manager/ Support Engineer	Suspend and Resume Data Pool Insert Actions	(P) 17.10.3	
4	Archive Manager/ Support Engineer	Check the Data Pool Insert Queue and Cancel a Data Pool Insert Action	(P) 17.10.4	
5	Archive Manager/ Support Engineer	Toggle the State of the NoFreeSpaceFlag	(P) 17.10.5	
6	Archive Manager/ Support Engineer	Configure the Number of Allowed Active Insert Processes	(P) 17.10.6	
7	Archive Manager/ Support Engineer	Configure the Default Retention Period and the Default Retention Priority	(P) 17.10.7	
8	Archive Manager/ Support Engineer	Use the DPM GUI to View Collection Groups and Collections	(P) 17.10.8	
9	Archive Manager/ Support Engineer	Use the DPM GUI to Modify Collection Groups	(P) 17.10.9	
10	Archive Manager/ Support Engineer	Use the DPM GUI to Add a Collection Group	(P) 17.10.10	

11	Archive Manager/ Support Engineer	Use the DPM GUI to Add Collections to a Collection Group	(P) 17.10.11	
12	Archive Manager/ Support Engineer	Use the DPM GUI to Modify Collections in a Collection Group	(P) 17.10.12	
13	Archive Manager/ Support Engineer	Use the Update Granule Utility to Extend the Retention for Selected Science Granules	(P) 17.10.13	
14	Archive Manager/ Support Engineer	Invoke the Data Pool Cleanup Utility Manually	(P) 17.10.14	
15	Archive Manager/ Support Engineer	Establish Data Pool Cleanup to Run with <i>cron</i>	(P) 17.10.15	
16	Archive Manager/ Support Engineer	Specify Data Pool Access Statistics Rollup Start Time and DPASU Execution with <i>cron</i>	(P) 17.10.16	
17	Archive Manager/ Support Engineer	Specify Data Pool Access Statistics Utility Execution from the Command Line	(P) 17.10.17	
18	Archive Manager/ Support Engineer	Archive Access Statistics using the Data Pool Archive Access Statistics Data Utility	(P) 17.10.18	
19	Archive Manager/ Support Engineer	Delete Access Statistics using the Data Pool Delete Access Statistics Data Utility	(P) 17.10.19	
20	Archive Manager/ Support Engineer	Restore Access Statistics using the Data Pool Restore Access Statistics Data Utility	(P) 17.10.20	

17.10.1 Launch the Data Pool Maintenance (DPM) GUI

The procedure for launching the GUI is provided separately here and referenced in other procedures. Table 17.10-2 presents the steps required to launch the DPM GUI. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 At the UNIX command shell prompt, type **setenv DISPLAY *clientname*:0.0** and then press the **Return/Enter** key.
 - For *clientname*, use either the local terminal/workstation IP address or its machine name.

- 2 Start the log-in to a Netscape host by typing **/tools/bin/ssh *hostname*** (e.g., g0ins02, e0ins02, l0ins02, n0ins02) at the UNIX command shell prompt, and press the **Return/Enter** key.
 - If you receive the message, **Host key not found from the list of known hosts. Are you sure you want to continue connecting (yes/no)?** type **yes** (“y” alone does not work).
 - If you have previously set up a secure shell passphrase and executed **sshremote**, a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears; continue with Step 3.
 - If you have not previously set up a secure shell passphrase; go to Step 4.
- 3 If a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears, type your **Passphrase** and then press the **Return/Enter** key. Go to Step 5.
- 4 At the **<user@remotehost>'s password:** prompt, type your **Password** and then press the **Return/Enter** key.
 - You are logged in and a UNIX command shell prompt is displayed.
- 5 Type **netscape** and then press the **Return/Enter** key.
 - The Netscape web browser is displayed.
- 6 Click in the **Netsite:** field.
 - The field is highlighted.
- 7 Type the Universal Resource Locator (URL) for the DPM GUI and then press the **Return/Enter** key.
 - The DPM **Home Page** is displayed, offering access to Data Pool maintenance functions (**List Insert Queue**, **Manage Configuration Parameters**, and **Manage Collection Groups**) as well as a tab for **Help** in navigating the GUI.

Table 17.10-2. Launch the DPM GUI

Step	What to Do	Action to Take
1	setenv DISPLAY <i>clientname</i>:0.0	enter text; press Return/Enter
2	/tools/bin/ssh <i>hostname</i>	enter text; press Return/Enter
3	Passphrase (or Step 4)	enter text; press Return/Enter
4	Password	enter text; press Return/Enter
5	netscape	enter text; press Return/Enter
6	Move cursor to Netsite: field	single-click
7	http://<URL>	enter text; press Return/Enter

17.10.2 Using the DPM GUI to Monitor Data Pool Active Insert Processes

It may be useful to keep an instance of the DPM GUI displayed to monitor Data Pool Active Insert Processes. Table 17.10-3 presents the steps required to use the DPM GUI to monitor Data Pool active insert processes. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DPM GUI (refer to Procedure 17.10.1 **Launch the DPM GUI**).
 - The **Home Page** is the default display, providing a table of active insert processes showing columns of detailed information for each process, including:
 - the UNIX process identifier (**Unix ProcessId**).
 - the database identifier (**dbID**), or Granule ID for the granule being processed.
 - the **Collection** to which the granule belongs.
 - the **Version** for the collection to which the granule belongs.
 - the time at which the insert processing started (**StartTime**).
 - the time at which the status listed in the next column was achieved (**StatusTime**).
 - the current state of the insert process (**Status**).
 - the **AMASS Cache** availability (**Y** or **N**) of the granule being processed (**Note**: The system is designed for rapid insertion of data into the Data Pool by quickly processing data that are available in cache, such as data that are staged for archiving. If the insert processing is delayed and the data are removed from cache, the Data Pool insert is likely to fail.)
 - the number of attempts (**Retries**) by the process to recover from retrievable errors (e.g., Data Pool disk temporarily unavailable, Data Pool directory does not exist, Data Pool database temporarily unavailable).
- 2 If it is desired to obtain an immediate screen refresh, click on the **Refresh** button near the upper right corner of the display to refresh the screen.
 - The displayed insert action data are updated. (**Note**: The screen refreshes automatically at intervals determined by the number of seconds specified in the **Screen Refresh Rate** field.
- 3 If it is desired to change the automatic screen refresh rate, click at the end of the **Screen Refresh Rate** field.
 - The cursor is displayed at the end of the **Screen Refresh Rate** field.

- 4 To set a new value for **Screen Refresh Rate**, use the **Backspace** key to remove the current value, type the new value (in seconds), and then click on the **Apply** button to the right of the **Screen Refresh Rate** field.
 - The typed value appears in the **Screen Refresh Rate** field and the automatic refresh frequency is based on the new value.

Table 17.10-3. Using the DPM GUI to Monitor Data Pool Active Insert Processes

Step	What to Do	Action to Take
1	Launch the DPM GUI	use Procedure 17.10.1
2	If desired, activate the Refresh button	single-click
3	If desired, move cursor to Screen Refresh Rate field	single-click
4	Enter new value and activate the Apply button	enter text; single-click

17.10.3 Suspend and Resume Data Pool Insert Actions

At times, it may be necessary to suspend Data Pool insert actions (e.g., for cleanup or other maintenance activities). Table 17.10-4 presents the steps required to suspend and resume Data Pool insert actions. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DPM GUI (refer to Procedure 17.10.1 **Launch the DPM GUI**).
 - The **Home Page** is the default display, providing status and capabilities for changing the **Screen Refresh Rate**, **NoFreeSpaceFlag**, and **Suspension** state, as well as a table of active insert processes.
- 2 To suspend Data Pool insert actions, click the **SUSPEND** selection/indicator button.
 - The selected button is filled to indicate its selection.
- 3 Click on the **Apply** button to the right of the **SUSPEND** and **RESUME** selection/indicator buttons.
 - The screen is refreshed and the **Suspension** status line shows "Data Pool Insert actions were suspended on <date time>" to indicate the change.
- 4 To resume Data Pool insert actions, click the **RESUME** selection/indicator button.
 - The selected button is filled to indicate its selection.
- 5 Click on the **Apply** button to the right of the **SUSPEND** and **RESUME** selection/indicator buttons.
 - The screen is refreshed and the **Suspension** status line shows "Data Pool Insert actions were resumed on <date time>" to indicate the change.

Table 17.10-4. Suspend and Resume Data Pool Insert Actions

Step	What to Do	Action to Take
1	Launch the DPM GUI	use Procedure 17.10.1
2	Select SUSPEND	single-click
3	Activate the Apply button	single-click
4	Select RESUME	single-click
5	Activate the Apply button	single-click

17.10.4 Check the Data Pool Insert Queue and Cancel a Data Pool Insert Action

The **List Insert Queue** tab of the DPM GUI provides a list of Data Pool inserts left to process. It also provides for each listed insert a check box permitting the insert to be marked for cancellation, and an **Apply Change** button to implement the cancellation. Table 17.10-5 presents the steps required to check the Data Pool insert queue and cancel a Data Pool insert action. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DPM GUI (refer to Procedure 17.10.1 **Launch the DPM GUI**).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**List Insert Queue**, **Manage Configuration Parameters**, and **Manage Collection Groups**).
- 2 Click on the **List Insert Queue** tab.
 - The **List Insert Queue** page is displayed, providing a table of inserts left to process showing columns of detailed information for each process, including:
 - the insert queue identifier (**InsertQueueID**).
 - the subscription identifier (**SubID**) of the subscription selected by the software for processing (*Note*: There may be multiple subscriptions specifying insertion of specific data into the Data Pool, but only one insert is needed and therefore only one of the subscriptions serves as the basis for the insert action. The **SubID** is of no particular significance to an operator and may safely be ignored.)
 - the database identifier (**dbID**), or Granule ID for the granule to be processed.
 - the **Collection** to which the granule belongs.
 - the **Version** for the collection to which the granule belongs.
 - an indication of whether the insert is to include science granules and metadata or just the metadata (**Science Granules and/or Metadata**).
 - the time at which the insert was placed in the insert queue (**Enqueue Time**).

- the number of attempts (**Retries**) by the process to recover from retrieable errors (e.g., Data Pool disk temporarily unavailable, Data Pool directory does not exist, Data Pool database temporarily unavailable).
 - There is a **Continue** link at the bottom of the page; if there are more inserts than can be displayed in the space of one page, this link permits display of the continuation of the list.
- 3 If it is desired to cancel one or more inserts, click on the check box at the end of the row of information for the insert(s) to be canceled.
- The check box for each selected insert is filled to indicate selection.
- 4 To implement the cancellation of any selected insert(s), click on the **Apply Change** button.
- A confirmation message asks "Are you ready to cancel the insert for . . ." and there are links displayed for **Yes, cancel insert** and **No, return to previous page**.
- 5 To confirm cancellation, click on the **Yes, cancel insert** link.
- The **List Insert Queue** page is displayed with the canceled insert(s) removed and the count of inserts left to process reduced by the number of inserts canceled.

Table 17.10-5. Check the Data Pool Insert Queue and Cancel a Data Pool Insert Action

Step	What to Do	Action to Take
1	Launch the DPM GUI	use Procedure 17.10.1
2	Select List Insert Queue tab	single-click
3	Mark selection box for any insert(s) to be canceled	single-click(s)
4	Activate the Apply Change button	single-click
5	Activate the Yes, cancel insert link	single-click

17.10.5 Toggle the State of the NoFreeSpaceFlag

The **NoFreeSpaceFlag** must be **OFF** in order to make inserts into the Data Pool. When an insert requires more space than is available on the Data Pool disks, the insert fails and the **NoFreeSpaceFlag** is set to **ON**. To enable further inserts, it is necessary to free up additional space (e.g., run the Data Pool Cleanup utility; see Procedure 17.10.14 **Invoke the Data Pool Cleanup Utility Manually**) and then toggle the **NoFreeSpaceFlag** back to **OFF**. Table 17.10-6 presents the steps required to toggle the state of the NoFreeSpaceFlag. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DPM GUI (refer to Procedure 17.10.1 **Launch the DPM GUI**).
 - The **Home Page** is the default display, providing status and capabilities for changing the **Screen Refresh Rate**, **NoFreeSpaceFlag**, and **Suspension** state, as well as a table of active insert processes.
- 2 To toggle the state of the **NoFreeSpaceFlag** from **ON** to **OFF** (i.e., if the **ON** selection/indicator is filled), click the **OFF** selection/indicator button.
 - The **OFF** selection/indicator button is filled to indicate its selection.
- 3 Click on the **Apply** button to the right of the **ON** and **OFF** selection/indicator buttons.
 - The screen is refreshed and the **NoFreeSpaceFlag** status line shows "NoFreeSpaceFlag was turned off <date/time>" to indicate the change.

Table 17.10-6. Toggle the State of the NoFreeSpaceFlag

Step	What to Do	Action to Take
1	Launch the DPM GUI	use Procedure 17.10.1
2	Select OFF	single-click
3	Activate the Apply button	single-click

17.10.6 Configure the Number of Allowed Active Insert Processes

The **Manage Configuration Parameters** tab on the DPM GUI permits setting or changing configuration parameters for a number of configuration defaults for the GUI. As noted in prior sections, some parameters can also be changed on the **Home Page** tab. Although most of the parameters managed on the **Manage Configuration Parameters** tab are not likely to be changed frequently, the operator may want to change some of them for tuning the Data Pool. Three of them are addressed in this section. The change process is essentially the same for any of the parameters listed on the tab.

Table 17.10-7 presents the steps required to configure the number of allowed active insert processes. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DPM GUI (refer to Procedure 17.10.1 **Launch the DPM GUI**).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**List Insert Queue**, **Manage Configuration Parameters**, and **Manage Collection Groups**).
- 2 Click on the **Manage Configuration Parameters** tab.
 - The **Manage Configuration Parameters** page is displayed, providing a table of parameters showing three columns: **Parameter Name**, **Parameter Value** (including an entry field with current value, followed by a brief description of the parameter),

and **Click on Box to Modify Parm** (containing a check box to mark the parameter for change).

- There is an **Apply Change** button at the bottom of the page to implement any selected change(s).
- 3 In the row for the **NumOfAllowedInsertProcesses** parameter, click at the end of the entry field in the **Parameter Value** column.
 - The cursor is displayed at the end of the entry field.
 - 4 Use the **Backspace** key to remove the current value, and then type the desired value.
 - The typed entry is displayed in the field.
 - 5 In the row for the **NumOfAllowedInsertProcesses** parameter, click in the check box in the **Click on Box to Modify Parm** column.
 - The box is filled to indicate selection.
 - 6 Click the **Apply Change** button.
 - The screen is refreshed, the check box is unfilled, and the displayed **Parameter Value** reflects the change.

Table 17.10-7. Configure the Number of Allowed Active Insert Processes

Step	What to Do	Action to Take
1	Launch the DPM GUI	use Procedure 17.10.1
2	Select Manage Configuration Parameters tab	single-click
3	In the row for NumOfAllowedInsertProcesses , move cursor to end of Parameter Value entry field	single-click
4	Backspace to remove current value and type the new value	press Backspace key; enter text
5	In the row for NumOfAllowedInsertProcesses , mark check box in Click on Box to Modify Parm column	single-click
6	Activate the Apply Change button	single-click

17.10.7 Configure the Default Retention Period and the Default Retention Priority

Table 17.10-8 presents the steps required to configure the default retention period and the default retention priority. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DPM GUI (refer to Procedure 17.10.1 **Launch the DPM GUI**).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**List Insert Queue**, **Manage Configuration Parameters**, and **Manage Collection Groups**).
- 2 Click on the **Manage Configuration Parameters** tab.
 - The **Manage Configuration Parameters** page is displayed, providing a table of parameters showing three columns: **Parameter Name**, **Parameter Value** (including an entry field with current value, followed by a brief description of the parameter), and **Click on Box to Modify Parm** (containing a check box to mark the parameter for change).
 - There is an **Apply Change** button at the bottom of the page to implement any selected change(s).
- 3 In the row for the **DefaultRetentionPeriod** parameter, click at the end of the entry field in the **Parameter Value** column.
 - The cursor is displayed at the end of the entry field.
- 4 Use the **Backspace** key to remove the current value, and then type the desired value.
 - The typed entry is displayed in the field.
- 5 In the row for the **DefaultRetentionPeriod** parameter, click in the check box in the **Click on Box to Modify Parm** column.
 - The box is filled to indicate selection.
- 6 In the row for the **DefaultRetentionPriority** parameter, click at the end of the entry field in the **Parameter Value** column.
 - The cursor is displayed at the end of the entry field.
- 7 Use the **Backspace** key to remove the current value, and then type the desired value.
 - The typed entry is displayed in the field.
- 8 In the row for the **DefaultRetentionPriority** parameter, click in the check box in the **Click on Box to Modify Parm** column.
 - The box is filled to indicate selection.
- 9 Click the **Apply Change** button.
 - The screen is refreshed, the check boxes are unfilled, and the displayed **Parameter Value** reflects the changes.

Table 17.10-8. Configure the Default Retention Period and the Default Retention Priority

Step	What to Do	Action to Take
1	Launch the DPM GUI	use Procedure 17.10.1
2	Select Manage Configuration Parameters tab	single-click
3	In the row for DefaultRetentionPeriod , move cursor to end of Parameter Value entry field	single-click
4	Backspace to remove current value and type the new value	press Backspace key; enter text
5	In the row for DefaultRetentionPeriod , mark check box in Click on Box to Modify Parm column	single-click
6	In the row for DefaultRetentionPriority , move cursor to end of Parameter Value entry field	single-click
7	Backspace to remove current value and type the new value	press Backspace key; enter text
8	In the row for DefaultRetentionPriority , mark check box in Click on Box to Modify Parm column	single-click
9	Activate the Apply Change button	single-click

17.10.8 Use the DPM GUI to View Collection Groups and Collections

The **Manage Collection Groups** tab of the DPM GUI allows an operator to view and modify collection groups or to add a collection group in the Data Pool database. It also provides access to screens for viewing and modifying collections within a collection group, or for adding a collection to a collection group. Table 17.10-9 presents the steps required to use the DPM GUI to view collection groups and collections. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DPM GUI (refer to Procedure 17.10.1 **Launch the DPM GUI**).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**List Insert Queue**, **Manage Configuration Parameters**, and **Manage Collection Groups**).
- 2 Click on the **Manage Collection Groups** tab.
 - The **Manage Collection Groups** page is displayed, providing a table listing collection groups with columns providing for each group a **Group ID** and a brief **Description** of the group. The entries in the **Group ID** column are links.
 - At the bottom of the page there are links permitting access to functions for **Add Collection Group** and **Modify Collection Group**.

- 3 To obtain more information about the collections in one of the groups, click on its link in the **Group ID** column.
 - The **List of Collections** page is displayed, identifying the selected group at the top and listing for that group the collections that are valid for insertion in the Data Pool, as well as whether the insertion is for science and metadata or metadata only. On this page, the collection ID entries in the **Collection** column are links.
- 4 To obtain a description of one of the collections, click on its link in the **Collection** column.
 - A **Description for Collection** page is displayed identifying the selected collection and providing the description. There is a link at the bottom to enable the operator to **Return to previous page**.
- 5 If it is desirable to view a description for another collection in the same group, click on the **Return to previous page** link; otherwise, go to Step 7.
 - The **List of Collections** page is displayed again.
- 6 Repeat Steps 4 and 5 as desired.
- 7 If it is desired to obtain more information about collections in a different collection group, repeat Steps 2 through 6.

Table 17.10-9. Use the DPM GUI to View Collection Groups and Collections

Step	What to Do	Action to Take
1	Launch the DPM GUI	use Procedure 17.10.1
2	Select Manage Collection Groups tab	single-click
3	In the Group ID column, activate the link for a collection for which more information is to be displayed	single-click
4	In the Collection column, activate the link for a collection for which a description is to be displayed	single-click
5	To view a description for another collection in the same group, activate the Return to previous page link; otherwise, go to Step 7	single-click
6	Repeat Steps 4 and 5 as desired	
7	To obtain more information about collections in a different group, repeat Steps 2 through 6	

17.10.9 Use the DPM GUI to Modify Collection Groups

Rarely, it may be desirable to modify the description of one or more of the collection groups listed on the **Manage Collection Groups** page. Table 17.10-10 presents the steps required to use the DPM GUI to modify collection groups. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DPM GUI (refer to Procedure 17.10.1 **Launch the DPM GUI**).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**List Insert Queue**, **Manage Configuration Parameters**, and **Manage Collection Groups**).
- 2 Click on the **Manage Collection Groups** tab.
 - The **Manage Collection Groups** page is displayed, providing a table listing collection groups with columns providing for each group a **Group ID** and a brief **Description** of the group. The entries in the **Group ID** column are links.
 - At the bottom of the page there are links permitting access to functions for **Add Collection Group** and **Modify Collection Group**.
- 3 Click on the **Modify Collection Group** link at the bottom of the page.
 - The screen displays a page listing the collection groups with their descriptions in text-entry fields, each accompanied by a check box in a **Click on Box to Modify Desc** column.
- 4 Click at the end of the text-entry field for the collection group to be modified.
 - The cursor is displayed at the end of the **Description** text-entry field.
- 5 Use the **Backspace** key to delete the existing description and type the desired description (*Note:* You may also use the mouse to drag the cursor over all or part of the text to be replaced, highlighting the selected text, and type new text to replace the highlighted text.)
 - The typed text appears in the field.
- 6 Click on the check box in the **Click on Box to Modify Desc** column to the right of the newly modified description.
 - The box is filled to indicate its selection.

- 7 Repeat Steps 4 through 6 for any additional descriptions to be modified.
- 8 Click on the **Apply Change** button at the bottom of the **Click on Box to Modify Desc** column.
 - The screen is refreshed and the **Manage Collection Groups** page reflects the changed description(s).

Table 17.10-10. Use the DPM GUI to Modify Collection Groups

Step	What to Do	Action to Take
1	Launch the DPM GUI	use Procedure 17.10.1
2	Select Manage Collection Groups tab	single-click
3	Activate the Modify Collection Group link	single-click
4	Move cursor to end of the text-entry field for the collection group to be modified	single-click
5	Backspace to remove current value and type the new description	press Backspace key; enter text
6	Mark the check box in the Click on Box to Modify Desc column to the right of the newly modified description	single-click
7	Repeat Steps 4 through 6 for any additional descriptions to be modified	
8	Activate the Apply Change button (at the bottom of the Click on Box to Modify Desc column)	single-click

17.10.10 Use the DPM GUI to Add a Collection Group

From time to time, it may be necessary to add a collection group (e.g., if a DAAC begins archiving data from a new instrument). Table 17.10-11 presents the steps required to use the DPM GUI to add a collection group. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DPM GUI (refer to Procedure 17.10.1 **Launch the DPM GUI**).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**List Insert Queue**, **Manage Configuration Parameters**, and **Manage Collection Groups**).

- 2 Click on the **Manage Collection Groups** tab.
 - The **Manage Collection Groups** page is displayed, providing a table listing collection groups with columns providing for each group a **Group ID** and a brief **Description** of the group. The entries in the **Group ID** column are links.
 - At the bottom of the page there are links permitting access to functions for **Add Collection Group** and **Modify Collection Group**.
- 3 Click on the **Add Collection Group** link at the bottom of the page.
 - The screen displays a page with columns of text-entry fields; a **Group ID** column and a **Description** column permit identification and description of new collection groups.
- 4 Click in a **Group ID** text-entry field.
 - The cursor is displayed in the **Group ID** text-entry field.
- 5 Type the identifier for a new collection group.
 - The typed entry is displayed in the field.
- 6 Click in the **Description** text-entry field for the newly entered **Group ID**.
 - The cursor is displayed at the end of the **Description** text-entry field.
- 7 Type the description for the new collection group.
 - The typed entry is displayed in the field.
- 8 Repeat Steps 4 through 7 for any additional collection groups to be added.
- 9 Click on the **Apply Change** button at the bottom of the entry-field area.
 - The screen is refreshed and the **Manage Collection Groups** page reflects the new collection group(s).

Table 17.10-11. Use the DPM GUI to Add a Collection Group

Step	What to Do	Action to Take
1	Launch the DPM GUI	use Procedure 17.10.1
2	Select Manage Collection Groups tab	single-click
3	Activate the Add Collection Group link	single-click
4	Move cursor to a Group ID text-entry field	single-click
5	Type the identifier for a new collection group	enter text
6	Move cursor to the Description text-entry field for the newly entered Group ID	single-click
7	Type the description for the new collection group	enter text
8	Repeat Steps 4 through 7 for any additional collection groups to be added	
9	Activate the Apply Change button at the bottom of the entry-field area	single-click

17.10.11 Use the DPM GUI to Add Collections to a Collection Group

Although an initial Data Pool structure is provided, not all collections are necessarily specified as eligible for Data Pool insertion. Based on experience, or on changes in demand, a DAAC may wish to add one or more collections to a data group. Table 17.10-12 presents the steps required to use the DPM GUI to add collections to a collection group. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DPM GUI (refer to Procedure 17.10.1 **Launch the DPM GUI**).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**List Insert Queue**, **Manage Configuration Parameters**, and **Manage Collection Groups**).
- 2 Click on the **Manage Collection Groups** tab.
 - The **Manage Collection Groups** page is displayed, providing a table listing collection groups with columns providing for each group a **Group ID** and a brief **Description** of the group. The entries in the **Group ID** column are links.
- 3 Click on the **Group ID** link for the collection group to which the collection is to be added.
 - The **List of Collections** page is displayed, identifying the selected group at the top and listing for that group the collections that are valid for insertion in the Data Pool, as well as whether the insertion is for science and metadata or metadata only.
- 4 Click on the **Add Collection** link at the bottom.
 - The **Add Collections** page is displayed, indicating the selected collection group at the top and listing from the Science Data Server database collections not currently

approved as eligible for Data Pool insertion. The page also provides a column with option buttons permitting the operator to select **science and metadata** or **metadata only** and a column enabling selection of **valid for data pool** or **invalid for data pool**. The defaults are **science and metadata** and **invalid for data pool**, respectively. There is also a **Click on Box to Add Collection** column with check boxes to mark the collection(s) for addition, and an **Apply Change** button to implement the addition(s).

- 5 If the collection to be added is to permit insertion of **science and metadata**, go to Step 6; otherwise, click on its option button in the **Science Granules and/or Metadata** column and click to select **metadata only**.
 - The selected choice is displayed on the option button.
- 6 Click on the collection's option button in the **Data Pool Insertion** column and click to select **valid for data pool**.
 - The selected choice is displayed on the option button.
- 7 Click on the collection's check box in the **Click on Box to Add Collection** column.
 - The box is filled to indicate its selection.
- 8 Repeat Steps 5 through 7 for any additional collections to be added to the collection group identified at the top of the page.
- 9 Click on the **Apply Change** button at the bottom of the **Click on Box to Add Collection** column.
 - The **List of Collections** page reflects the addition(s).
- 10 Repeat Steps 2 through 9 for any additional collection(s) to be added to another collection group.

Table 17.10-12. Use the DPM GUI to Add Collections to a Collection Group

Step	What to Do	Action to Take
1	Launch the DPM GUI	use Procedure 17.10.1
2	Select Manage Collection Groups tab	single-click
3	Activate the Group ID link for the collection group to which the collection is to be added	single-click
4	Activate the Add Collection link	single-click
5	To restrict insertion of metadata only, use the option button in the Science Granules and/or Metadata column to select metadata only ; otherwise, go to Step 6	click option
6	Use option button in the Data Pool Insertion column to select valid for data pool	click option
7	Mark the check box for the collection in the Click on Box to Add Collection column	single-click
8	Repeat Steps 5 through 7 for any additional collections to be added to the collection group identified at the top of the page	
9	Activate the Apply Change button at the bottom of the Click on Box to Add Collection column	single-click
10	Repeat Steps 2 through 9 for any additional collection(s) to be added to another collection group	

17.10.12 Use the DPM GUI to Modify Collections in a Collection Group

As part of managing the Data Pool storage and retention of data, making adjustments based on experience and/or changes in demand, it may be desirable to modify a collection. The modification may mean specifying that metadata only may continue to be inserted and science granules may no longer be inserted, or declaring the collection no longer valid for data pool insertion at all. Table 17.10-13 presents the steps required to use the DPM GUI to modify collections in a collection group. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Launch the DPM GUI (refer to Procedure 17.10.1 **Launch the DPM GUI**).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**List Insert Queue**, **Manage Configuration Parameters**, and **Manage Collection Groups**).

- 2 Click on the **Manage Collection Groups** tab.
 - The **Manage Collection Groups** page is displayed, providing a table listing collection groups with columns providing for each group a **Group ID** and a brief **Description** of the group. The entries in the **Group ID** column are links.
- 3 Click on the **Group ID** link for the collection group containing the collection to be modified.
 - The **List of Collections** page is displayed, identifying the selected group at the top and listing for that group the collections that are valid for insertion in the Data Pool, as well as whether the insertion is for science and metadata or metadata only.
- 4 Click on the **Modify Collection** link at the bottom.
 - The **Modify Collections** page is displayed, indicating the selected collection group at the top and listing collections currently approved as eligible for Data Pool insertion. The page also provides a column with option buttons permitting the operator to select **science and metadata** or **metadata only** and a column enabling selection of **valid for data pool** or **invalid for data pool**.
- 5 If the modification to a collection is to change the selected option displayed on the option button for insertion of **science and metadata** (in the **Science Granules and/or Metadata** column), click on that option button for the collection and click to select the desired option from the displayed choices.
 - The selected choice is displayed on the option button.
- 6 If the modification is to change a collection's eligibility for insertion in the Data Pool, click on that collection's option button in the **Data Pool Insertion** column and click to select the desired option from the displayed choices.
 - The selected choice is displayed on the option button.
- 7 Click on the collection's check box in the **Click on Box to Modify Collection** column.
 - The box is filled to indicate its selection.
- 8 Repeat Steps 5 through 7 for any additional collections to be modified.
- 9 Click on the **Apply Change** button at the bottom of the **Click on Box to Modify Collection** column.
 - The **List of Collections** page reflects the change(s).
- 10 Repeat Steps 2 through 9 for any modifications to one or more collections in another collection group.

Table 17.10-13. Use the DPM GUI to Modify Collections in a Collection Group

Step	What to Do	Action to Take
1	Launch the DPM GUI	use Procedure 17.10.1
2	Select Manage Collection Groups tab	single-click
3	Activate the Group ID link for the collection group containing the collection to be modified	single-click
4	Activate the Modify Collection link	single-click
5	To change the selected option displayed on the option button for insertion of science and metadata (in the Science Granules and/or Metadata column), use the option button in that column to select the desired option	click option
6	To change a collection's eligibility for insertion in the Data Pool, use option button in the Data Pool Insertion column to select the desired option	click option
7	Mark the check box for the collection in the Click on Box to Modify Collection column	single-click
8	Repeat Steps 5 through 7 for any additional collections to be modified	
9	Activate the Apply Change button at the bottom of the Click on Box to Modify Collection column	single-click
10	Repeat Steps 2 through 9 for any modifications to one or more collections in another collection group	

17.10.13 Use the Update Granule Utility to Extend the Retention for Selected Science Granules

A change in user interest in data from a particular location may make it desirable to retain certain data already in the Data Pool for a longer period of time than originally specified. Data Pool maintenance personnel can run the Update Granule Utility to update the expiration date for selected science granules. This utility also permits modifying a granule's retention priority, which can affect how soon the Data Pool Cleanup Utility removes the granule from the Data Pool.

The Update Granule Utility permits updating granule information using a command-line interface. It may be run using one (and only one) of the following three options, which must be specified:

- noprompt:** suppressing prompts and detailed information display.
- verbose:** providing detailed operator information and confirmation prompts.
- nonverbose:** providing only confirmation of granule updates.

A single granule may be updated using manual input. Multiple granule updates can only be handled using an input file containing a list of granules to be updated. The input file must be structured as a list of granules to be processed, one per line. Each line contains a granule ID (reflecting the Sybase entry in the Data Pool database), an expiration date, and (optionally) a new retention priority, the value of which may be null (i.e., left blank). The fields are separated by a

single space. There should be no blank lines before the first or after the last granule in the list. The file contents should be similar to the following example.

```
GRANULE_ID_4832 EXP_DATE=2002/2/28 RETENTION=255
GRANULE_ID_4876 EXP_DATE=2002/2/28 RETENTION=200
GRANULE_ID_4883 EXP_DATE=2002/2/28 RETENTION=
GRANULE_ID_4937 EXP_DATE=2002/2/28
GRANULE_ID_4966 EXP_DATE=2002/2/28 RETENTION=255
```

The Update Granule Utility connects to the Data Pool database and calls Sybase stored procedures to perform the requested updates. Therefore, the utility runs only if the Data Pool database server is running and if the database is available. It also assumes the stored procedures are present. The Granule Update Utility may be run as a background process, with suppression of all warning/error messages and confirmation prompts if desired. When the utility is run, it writes information, any warnings, any errors, and messages to a log file about granules as they are updated.

Table 17.10-14 presents the steps required to use the Update Granule Utility to extend the retention for selected science granules. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in at the machine on which the Update Granule Utility is installed (e.g., e0dig06, g0dig06, l0dig06, n0dig06).
- 2 To change to the directory for starting the Update Granule Utility, type **cd /usr/ecs/<MODE>/CUSTOM/utilities** and then press the **Return/Enter** key.
 - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/utilities**.
- 3 At the UNIX prompt, enter the command to start the Update Granule Utility, in the form **EcDIUpdateGranule.pl <command line parameters>**. (*Note:* The first command-line parameter specified must be **<MODE>**, a valid, existing Data Pool mode [e.g., OPS, TS1, TS2]).
 - The following six permutations are valid command-line entries for initiating the Update Granule utility:
 - **EcDIUpdateGranule.pl <MODE> -verbose -file <filename>** (to update granules listed in an input file named **<filename>** while displaying detailed information to the operator, and asking confirmation of the update).
 - **EcDIUpdateGranule.pl <MODE> -nonverbose -file <filename>** (to update granules listed in an input file named **<filename>** while suppressing detailed information to the operator, but still asking confirmation of the update).
 - **EcDIUpdateGranule.pl <MODE> -verbose -grnid <granuleID> -exp <expiration date> [-ret <retention priority>]** (to update a granule identified by its **<granuleID>** with a new expiration date and, optionally, a new retention priority while displaying detailed information to the operator, and asking confirmation of the update).

- **EcDIUpdateGranule.pl** *<MODE>* **-nonverbose -grnid** *<granuleID>* **-exp** *<expiration date>* [**-ret** *<retention priority>*] (to update a granule identified by its *<granuleID>* with a new expiration date and, optionally, a new retention priority while suppressing detailed information to the operator, but still asking confirmation of the update).
- **EcDIUpdateGranule.pl** *<MODE>* **-noprompt -file** *<filename>* (to update granules listed in an input file named *<filename>* with no confirmation or verbose information displayed to the operator).
- **EcDIUpdateGranule.pl** *<MODE>* **-noprompt -grnid** *<granuleID>* **-exp** *<expiration date>* [**-ret** *<retention priority>*] (to update a granule identified by its *<granuleID>* with a new expiration date and, optionally, a new retention priority with no confirmation or verbose information displayed to the operator).
- The utility executes and displays a confirmation prompt similar to the following:

```
You are about to start updating granules.
-----
Total number of granules: 11
Total size of granules: 8.61339673772454 MB
Do you wish to continue processing the update? [y/n]y
```

4 Type **y** and then press the **Return/Enter** key.

- The utility completes execution and displays output similar to the following:
- ```
Update completed.
Please check the database to ensure proper completion.

Update took 2 seconds to complete

Gracefully exiting...
```
- To check the database, have the Database Administrator use *isql* commands on the Data Pool database host to query the DIGranuleExpirationPriority table. It may also be useful to examine the Update Granule Utility log file to determine whether there were any problems with the execution. To examine that log file, go to Steps 5 and 6.

**5** To change to the directory containing the Update Granule Utility log file and other log files, type **cd /usr/ecs/<MODE>/CUSTOM/logs**, and then press the **Return/Enter** key.

- The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/logs**.

6 To examine the Update Granule Utility log file, type **pg EcDlUpdateGranule.log** and then press the **Return/Enter** key.

- The first page of the log file is displayed; additional sequential pages can be displayed by pressing the **Return/Enter** key at the **:** prompt. It is also possible to search forward by typing **/<search item>**. For example, to search the log file for reference to one of the granules updated, type **/<granuleID>** and then press the **Return/Enter** key.
- Although this procedure is written for the **pg** command, any UNIX editor or visualizing command (e.g., **vi**, **view**, **more**, **tail**) can be used to review the log.
- The log entries have a time and date stamp; about the time that the update was executed, the log should show entries similar to the following:

```
2001/11/29 15:52:50.814:Update started...
```

```
2001/11/29 15:52:50.964:Granule 4871 updated
2001/11/29 15:52:51.083:Granule 4954 updated
2001/11/29 15:52:51.212:Granule 4955 updated
2001/11/29 15:52:51.346:Granule 4956 updated
2001/11/29 15:52:51.409:Granule 4957 updated
2001/11/29 15:52:51.688:Granule 4959 updated
2001/11/29 15:52:51.778:Granule 4961 updated
2001/11/29 15:52:51.998:Granule 4963 updated
2001/11/29 15:52:52.107:Granule 4963 updated
2001/11/29 15:52:52.394:Granule 4964 updated
2001/11/29 15:52:52.569:Granule 4966 updated
2001/11/29 15:52:52.590:Update ended.
```

```
2001/11/29 15:52:52.608:This update took approximately 2 seconds
```

- If the log indicates errors or warnings, it may be necessary to correct the condition identified in the entry (e.g., edit the data in the granule list in the input file) and run the utility again. Specific error entries depend on the error that occurred; examples of error entries in the log may be similar to the following:

```
4959 AST_04 1 0.03962299 Jul 30 2001 12:00AM Feb 2 1998
11:59PM 255 2
```

```
Warning: The new expiration date for the above granule is less than
or equal to today's date.
```

```
DATABASE ERROR:Server message number=120001 severity=16 state=1
line=33 server=f2acg01_srvr procedure=ProcSelectGrExpiration
text=ProcSelectGrExpiration: Requested granule id not in database.
```

```
2001/11/29 15:50:36.647:Sybase Lookup ==> ERRORS WERE FOUND WITH
GRANULE "4654". (It may not exist or contains the wrong format).
```

```
2001/11/29 15:50:36.663:
```

```
EcDlUpdateGranule_1.pl aborted due to insufficient processing data:
All the granule triplets had errors.
```

**Table 17.10-14. Use the Update Granule Utility to Extend the Retention for Selected Science Granules**

| Step | What to Do                                                               | Action to Take                        |
|------|--------------------------------------------------------------------------|---------------------------------------|
| 1    | Log in at host for Update Granule Utility                                | enter text; press <b>Return/Enter</b> |
| 2    | <b>cd /usr/ecs/&lt;MODE&gt;/CUSTOM/utilities</b>                         | enter text; press <b>Return/Enter</b> |
| 3    | <b>EcDIUpdateGranule.pl &lt;command line parameters&gt;</b>              | enter text; press <b>Return/Enter</b> |
| 4    | Enter <b>y</b>                                                           | enter text; press <b>Return/Enter</b> |
| 5    | <b>cd /usr/ecs/&lt;MODE&gt;/CUSTOM/logs</b>                              | enter text; press <b>Return/Enter</b> |
| 6    | <b>pg EcDIUpdateGranule.log</b> (or other editor or visualizing command) | enter text; press <b>Return/Enter</b> |

#### 17.10.14 Invoke the Data Pool Cleanup Utility Manually

The Data Pool Cleanup Utility permits ECS Operations Staff to remove expired granules from the Data Pool disks and corresponding inventory. It must be executed on the machine where the granules are located. Qualification for cleanup is based on two criteria: expiration date/time and retention priority.

To determine whether a granule qualifies for deletion, the utility first compares the granule's expiration date (insert date plus retention period in days specified in the insert subscription) with a cut-off date/time. If a granule's expiration date is prior to the cut-off, the granule qualifies as expired.

- The default cut-off date/time is set to midnight of the previous day.
- The operator is permitted to specify an 'offset' in hours to add or subtract hours from the previous midnight to determine a cut-off date/time for deletion.

Next, the utility compares the granule's retention priority with any priority limit the operator has specified to identify those granules that should be retained in the Data Pool even though their expiration date has passed.

- Retention priority is an integer from 1 to 255.
- Retention priority for granules already in the Data Pool may be modified using the granule expiration update script.

The Data Pool Cleanup Utility removes those granules with expiration date prior to the cut-off date/time and with retention priority is less than or equal to the specified limit. If a priority limit is not specified in command-line input parameters at the time it is invoked, the Cleanup Utility reads the parameter 'DEFAULT\_LIMIT' from its configuration file to get a priority limit. If the operator does not wish to use retention priority as a criterion for deletion, the default limit should be set to 255.

The Cleanup Utility can alternatively take as input a file listing granuleId's for granules to be deleted. The file can contain single or multiple granuleId's per line separated by whitespace.

If the Cleanup Utility is interrupted during execution, upon restart it continues from the point of interruption. Further, in the interest of low database contention, the Cleanup Utility allows only one instance of itself to execute. The Cleanup Utility also provides an option to suppress operator prompts/messages, although there are not many. Upon completion of removing the

granule files from the Data Pool disks, the Cleanup Utility determines if there is sufficient free space to update the 'NoFreeSpace' flag in the Data Pool database, if necessary.

The utility may be executed using a **-noprompt** argument to suppress all confirmations and warnings normally displayed to standard output. Table 17.10-15 presents the steps required to invoke the Data Pool Cleanup Utility manually. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in at the machine where the Data Pool Cleanup Utility is installed (e.g., e0dig06, g0dig06, l0dig06, n0dig06).
  - **Note:** The operator who is executing the script must have the privilege to remove science, metadata, and browse files from the Data Pool disks.
- 2 To change to the directory for starting the Data Pool Cleanup Utility, type **cd /usr/ecs/<MODE>/CUSTOM/utilities** and then press the **Return/Enter** key.
  - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/utilities**.
- 3 At the UNIX prompt, enter the command to start the Cleanup Utility, in the form **EcDlCleanupDataPool.pl <MODE> [ -option1 <value1> . . . -optionN <valueN> ]**, specifying up to three options.
  - The following 10 permutations are valid command-line entries for initiating the Data Pool Cleanup utility:
    - **EcDlCleanupDataPool.pl <MODE>** (to delete all granules with retention priority less than or equal to the configured default limit that have expiration dates before midnight of the previous day).
    - **EcDlCleanupDataPool.pl <MODE> -noprompt** (to delete granules with retention priority less than or equal to the configured default limit that have expiration dates before midnight of the previous day while suppressing all operator prompts and confirmations).
    - **EcDlCleanupDataPool.pl <MODE> -limit <priority limit>** (to delete all granules with retention priority less than or equal to the specified *<priority limit>* that have expiration dates before midnight of the previous day).
    - **EcDlCleanupDataPool.pl <MODE> -limit <priority limit> -noprompt** (to delete all granules with retention priority less than or equal to the specified *<priority limit>* that have expiration dates before midnight of the previous day while suppressing all operator prompts and confirmations).
    - **EcDlCleanupDataPool.pl <MODE> -offset <±hours>** (to delete granules with retention priority less than or equal to the default limit that have expiration dates before midnight of the previous day plus or minus the specified number of *<hours>* -- e.g., **EcDlCleanupDataPool.pl OPS -offset -5** deletes granules with expiration date before 7:00 p.m. yesterday).

- **EcDlCleanupDataPool.pl <MODE> -offset <±hours> -noprompt** (to delete granules with retention priority less than or equal to the default limit that have expiration dates before midnight of the previous day plus or minus the specified number of <hours> while suppressing all operator prompts and confirmations).
  - **EcDlCleanupDataPool.pl <MODE> -offset <±hours> -limit <priority limit>** (to delete granules with retention priority less than or equal to the specified <priority limit> that have expiration dates before midnight of the previous day plus or minus the specified number of <hours>).
  - **EcDlCleanupDataPool.pl <MODE> -offset <±hours> -limit <priority limit> -noprompt** (to delete granules with retention priority less than or equal to the specified <priority limit> that have expiration dates before midnight of the previous day plus or minus the specified number of <hours> while suppressing all operator prompts and confirmations).
  - **EcDlCleanupDataPool.pl <MODE> -file <filename>** (to delete granules listed by granuleID in the input file named <filename>, and any associated browse granules). *Note:* The **-file** option may not be used with the **-offset** option or **-limit** option.
  - **EcDlCleanupDataPool.pl <MODE> -file <filename> -noprompt** (to delete granules listed by granuleID in the input file named <filename>, and any associated browse granules while suppressing all operator prompts and confirmations).
- The Cleanup Utility runs and the Cleanup Utility log file **EcDlCleanup.log** records errors, warnings, and information about utility events.

**Table 17.10-15. Invoke the Data Pool Cleanup Utility Manually**

| Step | What to Do                                                                                                                                                                         | Action to Take                        |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
| 1    | Log in at host for Data Pool Cleanup Utility                                                                                                                                       | enter text; press <b>Return/Enter</b> |
| 2    | <b>cd /usr/ecs/&lt;MODE&gt;/CUSTOM/utilities</b>                                                                                                                                   | enter text; press <b>Return/Enter</b> |
| 3    | <b>EcDlCleanupDataPool.pl &lt;MODE&gt; [-option1 &lt;value1&gt; . . . -optionN &lt;valueN&gt;]</b><br>(Options: <b>-noprompt</b> , <b>-limit</b> , <b>-offset</b> , <b>-file</b> ) | enter text; press <b>Return/Enter</b> |

### 17.10.15 Establish Data Pool Cleanup to Run (at 1:00 am) with *cron*

The Data Pool Cleanup Utility may be run with *cron* to execute it on a daily basis at a consistent time of day. The procedure specified here provides an example of adding a line to a *crontab* file to execute the OPS mode Data Pool Cleanup at 1:00 a.m. every day. Table 17.10-16 presents the steps required to establish Data Pool Cleanup to run with *cron*. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in at an ECS platform using an account with privileges to remove science, metadata, and browse files from Data Pool disks.
- 2 To ensure that the **crontab** command launches the **vi** editor, type **setenv EDITOR vi** and then press the **Return/Enter** key.
  - It may be desirable to include this command in the operator's **.cshrc** file to set the **crontab** editor to **vi** as part of the environmental settings normally used routinely.
- 3 Type **crontab -e** and then press the **Return/Enter** key.
  - The contents of the file are displayed, and the cursor is displayed on the first character at the upper left corner of the file. *Note:* If the operator has no **crontab** file on the current platform, this command opens a new one for editing.
- 4 If necessary, use the down arrow key on the keyboard to move the cursor down to a blank line.
  - The cursor is displayed at the beginning of the selected line.
- 5 Type **i** to put the **vi** editor into the insert mode.
  - The **vi** editor is in the insert mode, but no feedback is provided.
- 6 Type **0 1 \* \* \* /usr/ecs/OPS/CUSTOM/utilities/EcDlDataPoolCleanup.pl OPS**.
  - The typed entry appears to the left of the cursor.
- 7 Press the **Esc** key.
  - The cursor moves one character to the left and the **vi** editor is in the command mode.
- 8 Type **:wq** and then press the **Return/Enter** key.
  - UNIX displays a message identifying the number of lines and characters in the **crontab** file (stored in the directory **/var/spool/cron/crontabs**) and then displays the UNIX prompt.



**Table 17.10-16. Establish Data Pool Cleanup to Run (at 1:00 am) with cron**

| Step | What to Do                                                                                                                | Action to Take                        |
|------|---------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
| 1    | Log in at an ECS host using an account with privileges to remove science, metadata, and browse files from Data Pool disks | enter text; press <b>Return/Enter</b> |
| 2    | <b>setenv EDITOR vi</b>                                                                                                   | enter text; press <b>Return/Enter</b> |
| 3    | <b>crontab -e</b>                                                                                                         | enter text; press <b>Return/Enter</b> |
| 4    | If necessary, use <b>down arrow key</b> to move cursor to a blank line                                                    | press arrow key on keyboard           |
| 5    | To put <b>vi editor</b> in <b>insert</b> mode, type <b>i</b>                                                              | enter text command                    |
| 6    | <b>0 1 * * * /usr/ecs/OPS/CUSTOM/utilities/EcDIDataPoolCleanup.pl OPS</b>                                                 | enter text                            |
| 7    | To put <b>vi editor</b> in <b>command</b> mode, press <b>Esc</b> key                                                      | press <b>Esc</b> key on keyboard      |
| 8    | Exit <b>vi editor</b> with <b>:wq</b>                                                                                     | enter text; press <b>Return/Enter</b> |

#### **17.10.16 Specify Data Pool Access Statistics Rollup Start Time (at 1:00 am) and DPASU Execution (at 2:00 am), OPS Mode, with cron**

The Data Pool Access Statistics Utility (DPASU) parses logs of the Data Pool Web Access service and the FTP access service and stores the results in tables in the Data Pool database. The DPASU is a command-line utility that permits an option of entering input parameters. It is intended to be run with *cron* to cover an arbitrary 24-hour period starting at a time specified as a configuration parameter in a configuration file. However, an operator may run the utility from the command line specifying a start date as an input parameter to cover a period other than the normal 24-hour period addressed by *cron* or to cover that normal period if *cron* failed to process the logs for that period.

There are two versions of the DPASU, one for each type of log processed. The script named **EcDIRollupWebLogs.pl** runs on the Data Pool Web Access server and processes its log; its configuration file is **EcDIRollupWebLogs.CFG**. The script named **EcDIRollupFtpLogs.pl** runs on a server with access to SYSLOG with FTP access entries; its configuration file is **EcDIRollupFtpLogs.CFG**. These scripts capture data on downloads from the Data Pool, including date and time of access, path and file name of the file, and size of the file. The captured data are written to a temporary "flat file" -- a tab-delimited text file -- stored in the directory `/<ECS_HOME>/<MODE>/CUSTOM/data/DPL/`. The flat file is then exported to Sybase and stored in a table. The DPASU calls Sybase stored procedures to generate a separate rollup table, removes the flat file, and enters a record in a separate table identifying which periods have been rolled up in order to prevent inadvertent reprocessing of that period.

To prevent potential table locking, *cron* runs of the DPASU scripts should be separated so that they are not both running concurrently (e.g., separate their start times by at least 20 minutes). Use the following procedure to specify a 1:00 a.m. start time for the rollup and add a line to the *crontab* files to run the DPASU for the OPS mode beginning at 2:00 a.m. every day with a 20-minute separation between the scripts.

Table 17.10-17 presents the steps required to specify Data Pool access statistics rollup start time and DPASU execution with *cron*. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in at the host for EcDIRollupWebLogs.pl and its configuration file (e.g., e0mss21, g0mss21, l0mss21, n0mss21).
- 2 To change to the directory containing the configuration file, type the command **cd /usr/ecs/OPS/CUSTOM/cfg** and then press the **Return/Enter** key.
  - The working directory is changed to **/usr/ecs/OPS/CUSTOM/cfg**.
- 3 To look at the Rollup Start Time specified in the configuration file, type **vi EcDIRollupWebLogs.CFG** and then press the **Return/Enter** key.
  - The contents of the file are displayed, and the last line of the file indicates the start time in format similar to the following:  

```
ROLLUP_START_TIME=3:00
```

  
and the cursor is displayed on the first character at the upper left corner of the file.
  - If the start time is correct, exit **vi** by typing **:q!** and pressing the **Return/Enter** key; then go to Step 10. Otherwise, to change the time, execute Steps 4 - 9.
- 4 Use the arrow keys on the keyboard to move the cursor down to the line specifying the **ROLLUP\_START\_TIME** and to move it to the right until it is located over the first character in the time value.
  - The cursor is moved to the start time location; the line should look similar to the following:  

```
ROLLUP_START_TIME=█:00
```
- 5 Type **x** to delete the number under the cursor.
  - The number is deleted; the line should look similar to the following.  

```
ROLLUP_START_TIME=█00
```
  - **Note:** If more characters in the time value are to be changed, you can type **x** repeatedly to delete additional characters. For this exercise, you need only delete one character.
- 6 Type **i** to put the **vi** editor into the insert mode.
  - The **vi** editor is in the insert mode, but no feedback is provided.
- 7 Type **1**.
  - The typed entry appears to the left of the cursor.
- 8 Press the **Esc** key.
  - The cursor moves one character to the left and the **vi** editor is in the command mode.

- 9 Type **ZZ** (be sure to use upper case).
  - The file is saved and the UNIX prompt is displayed.
- 10 To ensure that the **crontab** command launches the **vi** editor, type **setenv EDITOR vi** and then press the **Return/Enter** key.
  - It may be desirable to include this command in the operator's **.cshrc** file to set the **crontab** editor to **vi** as part of the environmental settings normally used routinely.
- 11 Type **crontab -e** and then press the **Return/Enter** key.
  - The contents of the file are displayed, and the cursor is displayed on the first character at the upper left corner of the file. *Note*: If the operator has no **crontab** file on the current platform, this command opens a new one for editing.
- 12 If necessary, use the down arrow key on the keyboard to move the cursor down to a blank line.
  - The cursor is displayed at the beginning of the selected line.
- 13 Type **i** to put the **vi** editor into the insert mode.
  - The **vi** editor is in the insert mode, but no feedback is provided.
- 14 Type **0 2 \* \* \* /usr/ecs/OPS/CUSTOM/utilities/EcDIRollupWebLogs.pl OPS -noprompt**.
  - The typed entry appears to the left of the cursor.
- 15 Press the **Esc** key.
  - The cursor moves one character to the left and the **vi** editor is in the command mode.
- 16 Type **:wq** and then press the **Return/Enter** key.
  - UNIX displays a message identifying the number of lines and characters in the **crontab** file (stored in the directory **/var/spool/cron/crontabs**) and then displays the UNIX prompt.
- 17 Log in at the host for EcDIRollupFtpLogs.pl and its configuration file (e.g., e0dig06, g0dig06, l0dig06, n0dig06).
- 18 To change to the directory containing the configuration file, type the command **cd /usr/ecs/OPS/CUSTOM/cfg** and then press the **Return/Enter** key.
  - The working directory is changed to **/usr/ecs/OPS/CUSTOM/cfg**.

- 19** To look at the Rollup Start Time specified in the configuration file, type **vi EcDIRollupFtpLogs.CFG** and then press the **Return/Enter** key.
- The contents of the file are displayed, and the last line of the file indicates the start time in format similar to the following:  

```
ROLLUP_START_TIME=3:00
```

  
and the cursor is displayed on the first character at the upper left corner of the file.
  - If the start time is correct, exit **vi** by typing **:q!** and pressing the **Return/Enter** key; then go to Step 21. Otherwise, to change the time, execute Step 20.
- 20** Repeat Steps 4-9 to change the time in **EcDIRollupFtpLogs.CFG**.
- 21** To ensure that the **crontab** command launches the **vi** editor, type **setenv EDITOR vi** and then press the **Return/Enter** key.
- It may be desirable to include this command in the operator's **.cshrc** file to set the **crontab** editor to **vi** as part of the environmental settings normally used routinely.
- 22** Type **crontab -e** and then press the **Return/Enter** key.
- The contents of the file are displayed, and the cursor is displayed on the first character at the upper left corner of the file. **Note:** If the operator has no **crontab** file on the current platform, this command opens a new one for editing.
- 23** If necessary, use the down arrow key on the keyboard to move the cursor down to a blank line.
- The cursor is displayed at the beginning of the selected line.
- 24** Type **i** to put the **vi** editor into the insert mode.
- The **vi** editor is in the insert mode, but no feedback is provided.
- 25** Type **20 2 \* \* \* /usr/ecs/OPS/CUSTOM/utilities/EcDIRollupFtpLogs.pl OPS -noprompt**.
- The typed entry appears to the left of the cursor.
- 26** Press the **Esc** key.
- The cursor moves one character to the left and the **vi** editor is in the command mode.
- 27** Type **:wq** and then press the **Return/Enter** key.
- UNIX displays a message identifying the number of lines and characters in the **crontab** file (stored in the directory **/var/spool/cron/crontabs**) and then displays the UNIX prompt.

**Table 17.10-17. Specify Data Pool Access Statistics Rollup Start Time (at 1:00 am) and DPASU Execution (at 2:00 am), OPS Mode, with cron**

| Step | What to Do                                                                                                                   | Action to Take                        |
|------|------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
| 1    | Log in at the host for EcDIRollupWebLogs.pl and its configuration file                                                       | enter text; press <b>Return/Enter</b> |
| 2    | <b>cd /usr/ecs/OPS/CUSTOM/cfg</b>                                                                                            | enter text; press <b>Return/Enter</b> |
| 3    | <b>vi EcDIRollupWebLogs.CFG</b>                                                                                              | enter text; press <b>Return/Enter</b> |
| 4    | Move cursor <b>down</b> to line for <b>ROLLUP_START_TIME</b> and <b>right</b> until it is over first character in time value | press keyboard arrow keys             |
| 5    | To delete the number under the cursor, type <b>x</b>                                                                         | enter text command                    |
| 6    | To put <b>vi editor</b> in <b>insert</b> mode, type <b>i</b>                                                                 | enter text command                    |
| 7    | <b>1</b> (for <b>ROLLUP_START_TIME</b> of 1:00 am)                                                                           | enter text                            |
| 8    | To put <b>vi editor</b> in <b>command</b> mode, press <b>Esc</b> key                                                         | press <b>Esc</b> key on keyboard      |
| 9    | <b>ZZ</b>                                                                                                                    | enter text command                    |
| 10   | <b>setenv EDITOR vi</b>                                                                                                      | enter text; press <b>Return/Enter</b> |
| 11   | <b>crontab -e</b>                                                                                                            | enter text; press <b>Return/Enter</b> |
| 12   | If necessary, use <b>down arrow key</b> to move cursor to a blank line                                                       | press arrow key on keyboard           |
| 13   | To put <b>vi editor</b> in <b>insert</b> mode, type <b>i</b>                                                                 | enter text command                    |
| 14   | <b>0 2 * * * /usr/ecs/OPS/CUSTOM/utilities/EcDIRollupWebLogs.pl OPS -noprompt</b>                                            | enter text                            |
| 15   | To put <b>vi editor</b> in <b>command</b> mode, press <b>Esc</b> key                                                         | press <b>Esc</b> key on keyboard      |
| 16   | Exit <b>vi editor</b> with <b>:wq</b>                                                                                        | enter text; press <b>Return/Enter</b> |
| 17   | Log in at the host for EcDIRollupFtpLogs.pl and its configuration file                                                       | enter text; press <b>Return/Enter</b> |
| 18   | <b>cd /usr/ecs/OPS/CUSTOM/cfg</b>                                                                                            | enter text; press <b>Return/Enter</b> |
| 19   | <b>vi EcDIRollupFtpLogs.CFG</b>                                                                                              | enter text; press <b>Return/Enter</b> |
| 20   | Repeat steps 4 through 9 to change the time in <b>EcDIRollupFtpLogs.CFG</b>                                                  |                                       |
| 21   | <b>setenv EDITOR vi</b>                                                                                                      | enter text; press <b>Return/Enter</b> |
| 22   | <b>crontab -e</b>                                                                                                            | enter text; press <b>Return/Enter</b> |
| 23   | If necessary, use <b>down arrow key</b> to move cursor to a blank line                                                       | press arrow key on keyboard           |
| 24   | To put <b>vi editor</b> in <b>insert</b> mode, type <b>i</b>                                                                 | enter text command                    |
| 25   | <b>0 2 * * * /usr/ecs/OPS/CUSTOM/utilities/EcDIRollupFtpLogs.pl OPS -noprompt</b>                                            | enter text                            |
| 26   | To put <b>vi editor</b> in <b>command</b> mode, press <b>Esc</b> key                                                         | press <b>Esc</b> key on keyboard      |
| 27   | Exit <b>vi editor</b> with <b>:wq</b>                                                                                        | enter text; press <b>Return/Enter</b> |

### 17.10.17 Specify Data Pool Access Statistics Utility Execution from the Command Line

Although the Data Pool Access Statistics Utility scripts are intended to be run with **cron**, if it is necessary to run them from the command line, it is possible to do so. For example, if **cron** fails to complete successfully for any reason, no entry is made into the record table to indicate that a period was processed. In that event, the statistics can be captured for the missing interval by running the utility manually.

There are seven command-line parameters for use with the utility scripts:

- The **<MODE>** parameter indicates the mode (must specify a valid directory path) in which the script is to run; it is mandatory, unlabeled, and must be the first parameter following the command.
- The **-noprompt** parameter optionally specifies suppression of output to the screen.
- The **-nodelete** parameter optionally prevents the flat file from being deleted upon completion of the run.
- The **-flatfile <path/file>** parameter optionally provides an alternative path/file name for the flat file produced by the parser (useful only with the **-nodelete** option).
- The **-ftp <path/file>** parameter optionally indicates an alternative ftp log path/file(s) to be used instead of the configured default path/file (for the **EcDIRollupFtpLogs.pl** script only). Wildcards may be used, but must be escaped (i.e., preceded with a \).
- The **-web <path/file>** parameter optionally indicates an alternative web log path/file(s) to be used instead of the configured default path/file (for the **EcDIRollupWebLogs.pl** script only). Wildcards may be used, but must be escaped (i.e., preceded with a \).
- The **-start <date>** parameter optionally indicates an alternative start date for the rollup period, using the format MM/DD, and may be used to process a previously uncovered period.

With the exception of the mandatory **<MODE>** parameter, which must appear first after the command, the other parameters may be used in various orders and combinations. For example, to run without screen prompts or information, starting from December 22, and to retain the flat file, the command for accumulating statistics on web access should be entered as follows:

**EcDIRollupWebLogs.pl OPS -noprompt -nodelete -start 12/22.**

To run with normal screen information display, starting from February 15, but using an alternative file with wildcards for the web log, the command should be similar to the following:

**EcDIRollupWebLogs.pl OPS - start 2/15 -web /usr/var/\\*.log.**

Table 17.10-18 presents the steps required to specify Data Pool Access Statistics Utility execution from the command line, with normal screen information display. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in at the host for EcDIRollupWebLogs.pl and its configuration file (e.g., e0mss21, g0mss21, l0mss21, n0mss21).
- 2 To change directory to the directory containing the script, type the command **cd /usr/ecs/<MODE>/CUSTOM/utilities** and then press the **Return/Enter** key.
  - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/utilities**.
- 3 Type **EcDIRollupWebLogs.pl <MODE>** and then press the **Return/Enter** key.
  - The utility runs and displays information to the screen as it executes, in form similar to the following:

A Synergy II/Data Pool product

```
|_/_/_/_/_/_/_/_/_/_/_/_/_/_/_/_|
| 0 | | / 0 | \ \ | | | | |
| | | | / | | \ \ | \ / |
|_/_/_/_/_/_/_/_/_/_/_/_/_/_/__|
```

Data Pool Access Statistics Utility

Connecting to database...

The DPASU will examine the logs for access entries between the following times:

|        | Month | Day | Hour | Minute |
|--------|-------|-----|------|--------|
| START: | 11    | 26  | 03   | 00     |
| END:   | 11    | 27  | 02   | 59     |

Checking for already covered rollup periods...

File list:

/usr/ecs/OPS/COTS/www/ns-home/www/logs/access

Processing Web logs...

No access entries found in any of the Web logs

Cleaning up table "DIWebAccessLog"...OK

Exporting flat file to Sybase...OK

No access data was available to roll up.

DPASU will skip this step.

Rollup successful!

Removing flat file...OK

Gracefully exiting...

- 4 Log in at the host for EcDlRollupFtpLogs.pl and its configuration file (e.g., e0dig06, g0dig06, l0dig06, n0dig06).
- 5 To change directory to the directory containing the script, type the command **cd /usr/ecs/<MODE>/CUSTOM/utilities** and then press the **Return/Enter** key.
  - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/utilities**.
- 6 Type **EcDlRollupFtpLogs.pl <MODE>** and then press the **Return/Enter** key.
  - The utility runs and displays information to the screen as it executes, in form similar to the following:

A Synergy II/Data Pool product

```
/\|_/_/_/_/_/_/_/_/_/_/_/_/_/_/_									
0	_/_/ 0	\ \							
/			/		\		\ \ /		
//_/_/_/_/_/_/_/_/_/_/_/_/_/_/									
```

Data Pool Access Statistics Utility

Connecting to database...

The DPASU will examine the logs for access entries between the following times:

|        | Month | Day | Hour | Minute |
|--------|-------|-----|------|--------|
| START: | 11    | 26  | 03   | 00     |
| END:   | 11    | 27  | 02   | 59     |

Checking for already covered rollup periods...

File list:

/var/adm/SYSLOG

Processing FTP logs...

No access entries found in any of the FTP logs

Cleaning up table "DlFtpAccessLog"...OK

Exporting flat file to Sybase...OK

No access data was available to roll up.

DPASU will skip this step.

Rollup successful!

Removing flat file...OK

Gracefully exiting...



**Table 17.10-18. Specify Data Pool Access Statistics Utility Execution from the Command Line**

| Step | What to Do                                                                                                                                                                                                                          | Action to Take                        |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
| 1    | Log in at the host for EcDIRollupWebLogs.pl and its configuration file                                                                                                                                                              | enter text; press <b>Return/Enter</b> |
| 2    | <b>cd /usr/ecs/&lt;MODE&gt;/CUSTOM/utilities</b>                                                                                                                                                                                    | enter text; press <b>Return/Enter</b> |
| 3    | <b>EcDIRollupWebLogs.pl &lt;MODE&gt;</b><br>(Results in normal information display; options include: <b>-noprompt</b> , <b>-nodelete</b> , <b>-flatfile &lt;path/file&gt;</b> , <b>-web &lt;path/file&gt; -start &lt;date&gt;</b> ) | enter text; press <b>Return/Enter</b> |
| 4    | Log in at the host for EcDIRollupFtpLogs.pl and its configuration file                                                                                                                                                              | enter text; press <b>Return/Enter</b> |
| 5    | <b>cd /usr/ecs/&lt;MODE&gt;/CUSTOM/utilities</b>                                                                                                                                                                                    | enter text; press <b>Return/Enter</b> |
| 6    | <b>EcDIRollupFtpLogs.pl &lt;MODE&gt;</b><br>(Results in normal information display; options include: <b>-noprompt</b> , <b>-nodelete</b> , <b>-flatfile &lt;path/file&gt;</b> , <b>-ftp &lt;path/file&gt; -start &lt;date&gt;</b> ) | enter text; press <b>Return/Enter</b> |

#### 17.10.18 Archive Access Statistics using the Data Pool Archive Access Statistics Data Utility

The three remaining utilities are shell scripts for archiving, deleting, and restoring information in database tables populated by the DPASU. The **Data Pool Archive Access Statistics Data Utility** is run from the command line as needed or desirable to connect to the Data Pool database and write granule access data for a specified time range from the DIGranuleAccess, DIGranuleSubscription, and DIAccessRollup tables to an ASCII file. Once this is done, the operator can run the **Data Pool Delete Access Statistics Data Utility** from the command line to delete the archived data from the Data Pool database. If it is desirable to restore deleted data to the database, the **Data Pool Restore Access Statistics Data Utility** can be run from the command line to restore the data.

Table 17.10-19 presents the steps required to archive access statistics using the Data Pool Archive Access Statistics Data Utility. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in at the host for the Data Pool database (e.g., e0acg11, g0acg01, l0acg02, n0acg01).
- 2 To change directory to the directory containing the Data Pool Archive Access Statistics Data Utility, type **cd /usr/ecs/<MODE>/CUSTOM/dbms/DPL** and then press the **Return/Enter** key.
  - The working directory is changed to **cd /usr/ecs/<MODE>/CUSTOM/dbms/DPL**.

- 3 Type **DIDbArchiveAccessStat** *<MODE>* *<STARTDATE>* *<STOPDATE>* *<ARCHIVEDIR>* *<USERNAME>* *<SERVER>* *<DBNAME>* and then press the **Return/Enter** key.
  - *Note:* *<MODE>* is the mode in which the utility is being executed (e.g., OPS, TS1, TS2). *<STARTDATE>* is the start date time range, in format *yyyymmdd*, for the data to be archived. *<STOPDATE>* is the stop date time range, in format *yyyymmdd*, for the data to be archived. *<ARCHIVEDIR>* is the absolute path where the generated ASCII files are to be stored. *<USERNAME>* is the Sybase login name. *<SERVER>* is the Sybase Server for the Data Pool database (e.g., e0acg11\_srvr, g0acg01\_srvr, l0acg02\_srvr, n0acg01\_srvr). *<DBNAME>* is the name of the Data Pool database (e.g., DataPool\_OPS).
  - The script displays a prompt for entry of the password for the Sybase login.
- 4 Type *<password>* and then press the **Return/Enter** key (*Note:* This may require input from the Database Administrator).
  - The script runs and the Archive Access Statistics Utility log file **DIDbArchiveAccessStat.log** records errors, warnings, and information about utility events.

**Table 17.10-19. Archive Access Statistics using the Data Pool Archive Access Statistics Data Utility**

| Step | What to Do                                                                                                                                                                                      | Action to Take                        |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
| 1    | Log in at the host for the Data Pool database                                                                                                                                                   | enter text; press <b>Return/Enter</b> |
| 2    | <b>cd /usr/ecs/&lt;MODE&gt;/CUSTOM/dbms/DPL</b>                                                                                                                                                 | enter text; press <b>Return/Enter</b> |
| 3    | <b>DIDbArchiveAccessStat</b> <i>&lt;MODE&gt;</i> <i>&lt;STARTDATE&gt;</i> <i>&lt;STOPDATE&gt;</i> <i>&lt;ARCHIVEDIR&gt;</i> <i>&lt;USERNAME&gt;</i> <i>&lt;SERVER&gt;</i> <i>&lt;DBNAME&gt;</i> | enter text; press <b>Return/Enter</b> |
| 4    | <i>&lt;password&gt;</i> (from Database Administrator)                                                                                                                                           | enter text; press <b>Return/Enter</b> |

### 17.10.19 Delete Access Statistics using the Data Pool Delete Access Statistics Data Utility

Table 17.10-20 presents the steps required to delete access statistics using the Data Pool Delete Access Statistics Data Utility. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in at the host for the Data Pool database (e.g., e0acg11, g0acg01, l0acg02, n0acg01).
- 2 To change directory to the directory containing the Data Pool Delete Access Statistics Data Utility, type **cd /usr/ecs/<MODE>/CUSTOM/ dbms/DPL** and then press the **Return/Enter** key.
  - The working directory is changed to **cd /usr/ecs/<MODE>/CUSTOM/ dbms/DPL**.

3 Type **DIDbDeleteAccessStat** *<MODE>* *<STARTDATE>* *<STOPDATE>* *<USERNAME>* *<SERVER>* *<DBNAME>* and then press the **Return/Enter** key.

- *Note:* *<MODE>* is the mode in which the utility is being executed (e.g., OPS, TS1, TS2). *<STARTDATE>* is the start date time range, in format *yyyymmdd*, for the data to be deleted. *<STOPDATE>* is the stop date time range, in format *yyyymmdd*, for the data to be deleted. *<USERNAME>* is the Sybase login name. *<SERVER>* is the Sybase Server for the Data Pool database (e.g., e0acg11\_srvr, g0acg01\_srvr, l0acg02\_srvr, n0acg01\_srvr). *<DBNAME>* is the name of the Data Pool database (e.g., DataPool\_OPS).

- The script displays a prompt for entry of the password for the Sybase login.

4 Type *<password>* and then press the **Return/Enter** key (*Note:* This may require input from the Database Administrator).

- The script runs and the Delete Access Statistics Utility log file **DIDbDeleteAccessStat.log** records errors, warnings, and information about utility events.

**Table 17.10-20. Delete Access Statistics using the Data Pool Delete Access Statistics Data Utility**

| Step | What to Do                                                                                                                                                                                     | Action to Take                        |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
| 1    | Log in at the host for the Data Pool database                                                                                                                                                  | enter text; press <b>Return/Enter</b> |
| 2    | <b>cd /usr/ecs/</b> <i>&lt;MODE&gt;</i> <b>/CUSTOM/dbms/DPL</b>                                                                                                                                | enter text; press <b>Return/Enter</b> |
| 3    | <b>DIDbDeleteAccessStat</b> <i>&lt;MODE&gt;</i> <i>&lt;STARTDATE&gt;</i> <i>&lt;STOPDATE&gt;</i> <i>&lt;ARCHIVEDIR&gt;</i> <i>&lt;USERNAME&gt;</i> <i>&lt;SERVER&gt;</i> <i>&lt;DBNAME&gt;</i> | enter text; press <b>Return/Enter</b> |
| 4    | <i>&lt;password&gt;</i> (from Database Administrator)                                                                                                                                          | enter text; press <b>Return/Enter</b> |

### 17.10.20 Restore Access Statistics using the Data Pool Restore Access Statistics Data Utility

Table 17.10-21 presents the steps required to restore access statistics using the Data Pool Restore Access Statistics Data Utility. If you are already familiar with the procedure, you may prefer to use this quick-step table. If you are new to the system, or have not performed this task recently, you should use the following detailed procedure:

- 1 Log in at the host for the Data Pool database (e.g., e0acg11, g0acg01, l0acg02, n0acg01).
- 2 To change directory to the directory containing the Data Pool Restore Access Statistics Data Utility, type **cd /usr/ecs/***<MODE>***/CUSTOM/ dbms/DPL** and then press the **Return/Enter** key.

- The working directory is changed to **cd /usr/ecs/***<MODE>***/CUSTOM/ dbms/DPL**.

- 3 Type **DIDbRestoreAccessStat** *<MODE>* *<STARTDATE>* *<STOPDATE>* *<ARCHIVEDIR>* *<USERNAME>* *<SERVER>* *<DBNAME>* and then press the **Return/Enter** key.
  - *Note:* *<MODE>* is the mode in which the utility is being executed (e.g., OPS, TS1, TS2). *<STARTDATE>* is the start date time range, in format *yyyymmdd*, for the data to be restored. *<STOPDATE>* is the stop date time range, in format *yyyymmdd*, for the data to be restored. *<ARCHIVEDIR>* is the absolute path of the storage location for the ASCII files containing the data to be restored. *<USERNAME>* is the Sybase login name. *<SERVER>* is the Sybase Server for the Data Pool database (e.g., e0acg11\_svr, g0acg01\_svr, l0acg02\_svr, n0acg01\_svr). *<DBNAME>* is the name of the Data Pool database (e.g., DataPool OPS).
  - The script displays a prompt for entry of the password for the Sybase login.
- 4 Type *<password>* and then press the **Return/Enter** key (*Note:* This may require input from the Database Administrator).
  - The script runs and the Archive Access Statistics Utility log file **DIDbRestoreAccessStat.log** records errors, warnings, and information about utility events.

**Table 17.10-21. Restore Access Statistics using the  
Data Pool Restore Access Statistics Data Utility**

| Step | What to Do                                                                                                                                                                                      | Action to Take                        |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
| 1    | Log in at the host for the Data Pool database                                                                                                                                                   | enter text; press <b>Return/Enter</b> |
| 2    | <b>cd /usr/ecs/</b> <i>&lt;MODE&gt;</i> <b>/CUSTOM/dbms/DPL</b>                                                                                                                                 | enter text; press <b>Return/Enter</b> |
| 3    | <b>DIDbRestoreAccessStat</b> <i>&lt;MODE&gt;</i> <i>&lt;STARTDATE&gt;</i> <i>&lt;STOPDATE&gt;</i> <i>&lt;ARCHIVEDIR&gt;</i> <i>&lt;USERNAME&gt;</i> <i>&lt;SERVER&gt;</i> <i>&lt;DBNAME&gt;</i> | enter text; press <b>Return/Enter</b> |
| 4    | <i>&lt;password&gt;</i> (from Database Administrator)                                                                                                                                           | enter text; press <b>Return/Enter</b> |